# ONTAP 9.18.1 REST API reference

## REST API reference

NetApp
July 17, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-restapi-9181/index.html on July 17, 2026. Always check docs.netapp.com for the latest.

# Table of Contents


ONTAP 9.18.1 REST API reference . . . 1
REST API reference . . . 2
Getting started with the ONTAP REST API . . . 2
Using the ONTAP REST API online documentation . . . 2
Features for all ONTAP APIs . . . 2
Application . . . 37
ONTAP REST API Application endpoints . . . 37
Retrieve applications . . . 53
Create an application . . . 103
Retrieve application components . . . 160
Retrieve application component snapshots . . . 177
Create an application component snapshot . . . 184
Delete an application component snapshot . . . 189
Retrieve a snapshot for a specific application component . . . 194
Restore an application component snapshot . . . 199
Retrieve an application component . . . 202
Manage application snapshots . . . 219
Delete an application and all associated data . . . 244
Retrieve an application . . . 249
Update application properties . . . 294
Manage application consistency groups . . . 328
Retrieve consistency group metrics . . . 748
Manage application consistency group snapshots . . . 762
Manage application containers . . . 816
Retrieve application templates . . . 848
Retrieve an application template . . . 883
Storage cluster metrics . . . 917
Cloud . . . 919
ONTAP REST API Cloud endpoints . . . 919
Manage cloud targets . . . 919
Cluster . . . 976
ONTAP REST API Cluster endpoints . . . 976
Manage clusters . . . 977
Retrieve cluster chassis . . . 1126
Retrieve cluster counter tables . . . 1144
Manage cluster firmware history . . . 1182
View and manage cluster jobs . . . 1193
Retrieve capacity pool licenses . . . 1208
Manage cluster license managers . . . 1223
Manage cluster licensing . . . 1238
Manage the BlueXP cloud service . . . 1288
Manage cluster mediators . . . 1294
View and manage MetroCluster configurations . . . 1359

Display MetroCluster diagnostics . . . 1409
Manage MetroCluster DR groups . . . 1438
View and update MetroCluster interconnects . . . 1473
Retrieve MetroCluster node configurations . . . 1502
Retrieve MetroCluster operations . . . 1525
Retrieve MetroCluster SVMs . . . 1538
Manage cluster nodes . . . 1551
Cluster NTP . . . 1727
Manage cluster NTP keys . . . 1728
Manage cluster NTP servers . . . 1746
Manage cluster peers . . . 1772
Manage cluster schedules . . . 1828
Retrieve environment sensors . . . 1862
Manage environment sensors . . . 1869
Manage cluster software . . . 1876
Manage cluster software downloads . . . 1943
Manage cluster software history . . . 1955
Manage cluster software packages . . . 1962
View and update cluster web configurations . . . 1977
View tags used for resources . . . 1992
Manage resource tags . . . 1999
NAS . . . 2021
ONTAP REST API NAS endpoints . . . 2021
Manage Active Directory accounts . . . 2022
Manage Active Directory accounts for an SVM . . . 2043
Manage Active Directory domain controllers . . . 2062
Manage NAS audit configurations . . . 2079
Manage S3 audit configurations . . . 2132
View CIFS connection information . . . 2172
View CIFS domain-related information . . . 2181
Manage CIFS domain-related information . . . 2197
Manage CIFS DC configuration . . . 2219
View group policy objects . . . 2292
Manage home directory search paths . . . 2323
Manage local groups . . . 2342
Manage local group members . . . 2369
Manage local users . . . 2393
View NetBIOS information . . . 2426
Manage CIFS services . . . 2437
View CIFS active file sessions . . . 2579
View CIFS active sessions . . . 2602
View and update shadow copy information . . . 2628
Manage CIFS shares . . . 2666
Manage share-level ACL . . . 2732
Manage CIFS UNIX symlink mapping . . . 2759

Perform bulk import for CIFS user groups . . . 2790
Manage CIFS user group privileges . . . 2812
View file security permissions . . . 2836
Manage file security permissions and audit policies . . . 2843
Manage FPolicy configuration . . . 2925
View and update FPolicy server configuration . . . 3009
Manage FPolicy engine configuration . . . 3034
Manage FPolicy event configuration . . . 3093
Manage FPolicy persistent store configuration . . . 3138
Manage SVM FPolicy configuration . . . 3161
View lock details . . . 3214
View NFS client information . . . 3240
View and update NFS client settings . . . 3247
View NFS connected clients . . . 3255
Manage NFS export policies . . . 3266
View and update Kerberos interfaces . . . 3352
Manage Kerberos realms . . . 3371
Manage NFS services . . . 3399
Manage NFS over TLS interface configuration details . . . 3556
View and create Vscan configuration . . . 3574
Manage Vscan configuration . . . 3617
Manage Vscan On-Access policies . . . 3657
Manage Vscan On-Demand policies . . . 3692
Manage Vscan scanner-pool configuration . . . 3725
NDMP . . . 3761
ONTAP REST API NDMP endpoints . . . 3761
Manage NDMP mode . . . 3762
Manage NDMP node configurations . . . 3768
Manage NDMP session configurations . . . 3781
Manage NDMP configuration of SVM . . . 3804
Generate the NDMP user password . . . 3816
NVMe . . . 3820
ONTAP REST API NVMe endpoints . . . 3820
View NVMe interfaces . . . 3821
Manage NVMe services . . . 3843
View NVMe subsystem controllers . . . 3950
Manage NVMe subsystem maps . . . 3981
Manage NVMe subsystems . . . 4016
Manage NVMe namespaces . . . 4143
Name-services . . . 4314
ONTAP REST API Name-services endpoints . . . 4314
Manage group membership cache settings . . . 4314
Manage host cache settings . . . 4330
Manage netgroup cache settings . . . 4347
Manage global cache settings . . . 4364

Manage UNIX group cache settings . . . 4371
Manage UNIX user cache settings . . . 4389
Manage DNS configurations . . . 4406
View host records . . . 4462
Manage LDAP server configurations . . . 4469
Manage LDAP schemas . . . 4544
Manage local mapping for hostnames . . . 4587
Manage name mappings for SVMs . . . 4612
Display netgroup file details . . . 4643
Manage NIS server configurations . . . 4651
Manage UNIX group configurations . . . 4679
Manage UNIX group users . . . 4705
Manage UNIX user configurations . . . 4725
Networking . . . 4749
ONTAP REST API Networking endpoints . . . 4749
Manage broadcast domains . . . 4751
Manage network Ethernet ports . . . 4782
Retrieve network Ethernet switch ports . . . 4883
Manage Ethernet switches . . . 4913
Retrieve FC fabrics . . . 4947
Manage FC network interfaces . . . 5002
Retrieve FC port information . . . 5120
Manage HTTP proxy configuration . . . 5190
Manage BGP peer groups . . . 5218
Manage network IP interfaces . . . 5259
Manage network IP routes . . . 5357
Manage network IP service policies . . . 5389
Manage network IP subnets . . . 5427
Manage network IPspaces . . . 5469
Object-store . . . 5488
Manage S3 buckets . . . 5488
Manage protocols S3 services . . . 5624
Manage S3 service bucket configuration . . . 5790
Manage S3 bucket lifecycle management rules . . . 5935
Manage S3 bucket snapshots . . . 5976
Manage S3 service group configuration . . . 6001
Manage S3 service policies . . . 6040
Manage S3 service users . . . 6083
SAN . . . 6123
ONTAP REST API SAN endpoints . . . 6123
View FC login information . . . 6126
Manage FC WWPN aliases . . . 6150
Manage FC services for SVMs . . . 6171
Retrieve historical performance metrics for the FC protocol service of an SVM for a specific time . . . 6233
Manage SAN igroups . . . 6239

Manage SAN initiators . . . 6427
Manage iSCSI credentials . . . 6438
Manage iSCSI services . . . 6477
View iSCSI sessions . . . 6547
Manage SAN LUN maps . . . 6570
Add, remove, or discover LUN map reporting nodes . . . 6608
Manage SAN portsets . . . 6627
Manage SAN vVol bindings . . . 6683
Manage LUNs . . . 6712
Manage LUN attributes . . . 6964
Retrieve historical performance metrics for a LUN . . . 6988
Retrieve historical performance metrics for a LUN for a specific time . . . 6997
SVM . . . 7003
ONTAP REST API SVM endpoints . . . 7003
Manage SVM migration . . . 7003
Manage SVM peer permissions . . . 7089
Manage SVM peer relationships . . . 7111
Manage SVMs . . . 7145
Retrieve clients . . . 7317
Retrieve directories . . . 7331
Retrieve files . . . 7352
Retrieve users . . . 7372
Manage web services security configuration . . . 7389
Security . . . 7400
ONTAP REST API Security endpoints . . . 7400
Manage security-related operations . . . 7400
Manage security-related accounts . . . 7427
Manage scoped user accounts . . . 7453
Manage the anti-ransomware package . . . 7471
Manage anti-ransomware auto enablement . . . 7482
View suspect files generated by anti-ransomware . . . 7488
Delete suspect files . . . 7493
Manage data-entropy statistics . . . 7498
View and update audit settings . . . 7511
Forward audit logs to syslog/splunk servers . . . 7517
View administrative audit logs . . . 7546
Manage data SVM account information . . . 7555
Manage LDAP server configuration . . . 7568
Manage NIS configuration . . . 7616
View and update OAuth 2.0 status . . . 7633
View and create OAuth 2.0 configurations . . . 7637
View security key managers . . . 8797
Manage storage volumes . . . 11713
Manage files and directories . . . 12395
Manage volume snapshots . . . 12506

View I/O activity for clients . . . 12577
Manage directories . . . 12591
View I/O activity for files . . . 12614
View I/O activity for users . . . 12631
Support . . . 12645
ONTAP REST API Support endpoints . . . 12645
Manage an automatic update feature . . . 12647
Retrieve automatic update configuration . . . 12657
Manage automatic update configuration . . . 12663
Retrieve status for all automatic updates . . . 12672
Manage automatic updates . . . 12684
Manage AutoSupport configuration . . . 12695
Manage AutoSupport messages . . . 12717
Manage cluster configuration backup information . . . 12747
Manage configuration backups . . . 12754
Manage core dumps . . . 12772
Manage EMS configuration . . . 12788
Manage application logs . . . 12797
Manage event destinations . . . 12803
Manage an event destination instance . . . 12832
Retrieve events . . . 12856
Create or retrieve event filters . . . 12865
Manage event filters . . . 12884
Create or retrieve filter rules . . . 12902
Manage event filter rules . . . 12917
Manage the EMS event catalog . . . 12931
View and create EMS role-based configurations . . . 12938
Manage EMS role-based configurations . . . 12950
Manage cluster-wide SNMP configuration . . . 12961
Create or retrieve SNMP traphosts . . . 12971
Manage individual SNMP traphosts . . . 12985
Retrieve and create SNMP users . . . 12995
Manage individual SNMP users . . . 13017
Legal notices . . . 13034
Copyright . . . 13034
Trademarks . . . 13034
Patents . . . 13034
Privacy policy . . . 13034
API license . . . 13034
Open source . . . 13034

# ONTAP 9.18.1 REST API reference

# REST API reference

## Getting started with the ONTAP REST API

ONTAP adds support for an expansive RESTful API. The documentation below provides information about the types of API calls available with ONTAP, as well as details about using each API endpoint. You can learn more about the ONTAP REST API in the ONTAP automation doc site: https://docs.netapp.com/us-en/ontap-automation/index.html.

### Using the ONTAP REST API online documentation

Each API method includes usage examples, as well as a model that displays all the required and optional properties supported by the method. Click the *Model* link, available with each API method, to see all the required and optional properties supported by each method.

### Features for all ONTAP APIs

#### Getting started with the ONTAP REST API

Overview

Let's review some key things about RESTful APIs and how they're implemented in ONTAP:

- REST API URLs identify the resources that you'll be working with, including clusters, SVMs, and storage.
- REST APIs use HTTP methods GET, POST, PATCH, DELETE, and OPTIONS to indicate their actions.
- REST APIs return common HTTP status codes to indicate the results of each call. Additional error details can be included in the results body.
- REST APIs request and response bodies are encoded using JSON.
- REST APIs support hyperlinking among resources using the Content-Type "application/hal+json".
- REST API requests will be rejected when the wait queue is full. This is an uncommon occurrence and is designed to prevent a DOS attack.
- GET calls on collections usually return only name and UUID by default. If you want to retrieve additional properties, you need to specify them using the "fields" query parameter.
- ONTAP supports query-based DELETE or PATCH for all collection endpoints.
- REST APIs follows RFC3986 Section 5.2.4 regarding relative reference resolution.

If you're already familiar with the ONTAPI API (also known as ZAPI), there are some similarities between ONTAP REST APIs and ONTAPI. For example:

- Both support the same transport and security mechanisms.
- Both paginate results based on either number of seconds or number of records.
- Both support filtering the returned records based on property values.
- Both support limiting the returned properties.
- Both support concurrent requests. If ONTAP temporarily can't handle additional calls, it will respond with an HTTP error status code of 429 or 503 (depending on the kind of limit hit) and an error message in the body explaining the limit encountered.

However, there are important differences between REST APIs and the ONTAP CLI and ONTAPI that you should understand as well:

- In many cases, ONTAP REST APIs use different names for fields and features.
- REST APIs do not expose infrequently used CLI parameters.
- REST APIs do not treat the cluster or nodes as an SVM (aka Vserver).
- REST GET APIs support specifying a maximum time before paginating results. However, the default time is 15 seconds for REST (instead of 90 seconds for ONTAPI).
- REST APIs are generally ordered by UUID or ID, so a rename operation using the PATCH method doesn't change the path keys.
- REST APIs use one or more of the following properties to identify a resource: "name", "uuid", "id".
- REST APIs often execute the equivalent of multiple CLI commands in a single request.
- REST API properties use underscores instead of hyphens between words.
- REST API dates are always in ISO-8601 format.
- REST API comparisons between enum values (for <, >, ranges, and `order_by`) are done alphabetically. (In CLI and ONTAPI, enum comparisons are done based on an internal value for the enum.)
- REST API field '<' queries exclude records where the specified field is not set. You can add "\|null" (eg: limit=<10\|null) to also return records where the specified field is not set.

Relative reference resolution

In accordance with RFC3986 Section 5.2.4, ONTAP REST API removes dot segments to resolve the path. The examples below describe this behavior.

**This request uses non-URL-encoded dots:**

```
curl -siku <username>:<password> -X GET
"https://mgmt_ip_address/api/storage/volumes/<volume_uuid>/files/.."
```

**The previous request is resolved to this:**

```
curl -siku <username>:<password> -X GET
"https://mgmt_ip_address/api/storage/volumes/<volume_uuid>"
```

In the previous example, the volume information is returned rather than the file. Users should be cautious when using non-URL-encoded ".." since it can cause accidental data loss when making a delete request.

This request is interpreted as a delete to the volume rather than the file because of `...`

```
curl -siku <username>:<password> -X DELETE
"https://<mgmt_ip_address>/api/storage/volumes/<volume_uuid>/files/.."
```

To avoid this issue, dots should be URL-encoded to `%2E`. The previous request should look like this:

```
curl -siku <username>:<password> -X DELETE
"https://<mgmt_ip_address>/api/storage/volumes/<volume_uuid>/files/%2E%2E"
```

### HAL linking

Hypertext Application Language (HAL)

ONTAP REST APIs use HAL as the mechanism to support Hypermedia as the Engine of Application State (HATEOAS). When an object or attribute is returned that identifies a specific resource, a HAL-encoded link is also returned so that you can easily discover resources and be able to obtain more details about the resource.

Example

```
"aggregate": {
    "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
    "name": "aggr0",
    "_links": {
        "self": {
            "href": "/api/storage/aggregates/19425837-f2fa-4a9f-8f01-
712f626c983c"
        }
    }
}
```

### Query parameters

Overview

The following is a list of all the globally supported query parameters. This list is intended as a quick reference for syntax purposes. The query parameters are described in more detail in other sections of this documentation. Note that multiple queries can be combined using an "&".

```
# Request specific fields
fields=<field>[,...]

# Don't error due to an unknown field in "fields=" (default is false)
ignore_unknown_fields=<true|false>

# Query fields by value. If the value contains query characters
(*|,!<>..), it must be quoted to avoid their special query meaning
<field>=<query value>

# Return the records array
return_records=<true|false>

# Timeout and return after the specified number of seconds
return_timeout=<0..120 seconds>

# The number of records to collect (or act on for query-based
PATCH/DELETE) before returning
max_records=<number of records>

# Request a customized sort ordering
order_by=<field [asc|desc]>[,...]
$orderBy=<field [asc|desc]>[,...]

# Pretty print JSON response bodies
pretty=<true|false>

# Continue after encountering a failure. Only applicable to query-based
PATCH and DELETE.
continue_on_failure=<true|false>

# Begin returning records starting at an offset from the first record
offset=<offset from first record>
```

**Operations on multiple records**

Overview

Although they are not documented as individual methods in the list of REST APIs, every API supporting POST, PATCH, or DELETE, also supports operations on multiple records in the collection. These collection based operations can be specified as either a list of records, allowing each record to be mutated individually, or as a query that applies the same change to multiple records in the collection.

Option 1: Record-based POST, PATCH, and DELETE

APIs supporting POST, PATCH, or DELETE support these operations on the collection itself by specifying a `records` array in the JSON body. Each entry in this array must match the schema of an individual record of

the API. Parsing failures are returned synchronously. All other validations are performed asynchronously, including those that are usually returned synchronously for asynchronous APIs.

Rolling back changes on failure

POST and PATCH make a best-effort attempt to perform the operation atomically by rolling back any changes made to prior records if an operation on a later record fails. This behavior can be overridden via the `continue_on_failure` query parameter. When this parameter is *true* (default: *false*), the job performs the operation on every input record regardless of prior failures. If any failures are encountered in `continue_on_failure` mode, they are reported as the result of the job, but the successfully mutated records are not rolled back. DELETE operations do not attempt to roll back failures and always operate in `continue_on_failure` mode.

Serial and parallel operations

Most APIs support performing operations in parallel, and record-based operations take advantage of this, operating on multiple of the provided records at once. If there is some interdependence between the provided records, such as a hierarchical set of objects, the `serial_records=true` query parameter can be used to force the operation to be performed in order. APIs that do not support parallel operations run serially, by default. From a job tracking perspective, there is no difference between these two modes other than the speed at which the operation completes.

Asynchronous job and results

Unless the records array is directly documented in the API as synchronous, every record-based operation is performed as a job and follows the usual rules for asynchronous jobs. In addition to the usual `/api/cluster/jobs` record of the job, these operations return a job results href link. The link refers back to the API on which the operation was performed and can be followed after the job is complete. Job results include the current state of the records and fields operated on by the job. If a POST operation was successful on a record, that record is included in job results. If a DELETE operation was successful, that record is not included in the job results because it is no longer part of the collection. The same concept applies to failures. If the POST of a record fails, the record is not included in the job results because it is not part of the collection. The results also include any errors or rollback errors encountered by the job. The job results link is also returned as the Location header from a record-based POST job, supporting the same pattern for determining any generated identifiers as a single record POST. Note that the job results link does not support all of the usual GET features. Requesting specific fields and querying a subset of the records is supported, but *max_records* and *return_timeout* as well as other such query parameters for controlling the iteration of the collection is are not supported.

Unique identifiers go in the body

For PATCH and DELETE, any unique identifiers not included in the URI must be included in the JSON body of each record. For example, to DELETE a single volume, the API is usually `/api/storage/volumes/{uuid}`. To delete multiple volumes via a `records` array, the API is `/api/storage/volumes`, and the UUID of each volume is included in the JSON body of each record. For sub-endpoints with unique identifiers embedded in the path, such as `/api/storage/volumes/{volume.uuid}/files/{path}`, the `volume.uuid` is included once in the URI and is not included in each record. The `path` is included with each record.

Examples

Creating two volumes

```
# The API:
POST /api/storage/volumes
```

```
# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/volumes' -H 'accept:
application/hal+json' -d '{ "records": [
    { "svm": { "name": "svm1" }, "name" : "vol1", "size": "1GB",
"aggregates": [ { "name": "aggr1" } ] },
    { "svm": { "name": "svm1" }, "name" : "vol2", "size": "1GB",
"aggregates": [ { "name": "aggr1" } ] } ] }'
```

```
# The response:
{
  "job": {
    "uuid": "cc979874-a441-11eb-a707-005056bbbc41",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/cc979874-a441-11eb-a707-005056bbbc41"
      },
      "results": {
        "href": "/api/storage/volumes?job_results_uuid=cc979874-a441-11eb-
a707-005056bbbc41"
      }
    }
  }
}
```

```
# GET the job results:
curl -X GET 'https://<mgmt-
ip>/api/storage/volumes?job_results_uuid=cc979874-a441-11eb-a707-
005056bbbc41' -H 'accept: application/hal+json'
```

```
# The response:
{
  "records": [
    {
      "uuid": "cc9915c8-a441-11eb-a707-005056bbbc41",
      "name": "vol1",
      "size": 1073741824,
      "aggregates": [
```

```
        {
          "name": "aggr1",
          "uuid": "a589e251-a096-44be-b5d3-67ccf92d27e7"
        }
      ],
      "svm": {
        "name": "svm1"
      },
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cc9915c8-a441-11eb-a707-
005056bbbc41"
        }
      }
    },
    {
      "uuid": "cc9c1eea-a441-11eb-a707-005056bbbc41",
      "name": "vol2",
      "size": 1073741824,
      "aggregates": [
        {
          "name": "aggr1",
          "uuid": "a589e251-a096-44be-b5d3-67ccf92d27e7"
        }
      ],
      "svm": {
        "name": "svm1"
      },
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cc9c1eea-a441-11eb-a707-
005056bbbc41"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/volumes?job_results_uuid=cc979874-a441-11eb-
a707-005056bbbc41"
    }
  }
}
```


Creating two LUNs with a failure and an additional failure during rollback

Note that the LUN that was not rolled back is reported as being present in the results because it now exists in the collection.

```
# The API:
POST /api/storage/luns
```

```
# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/luns' -H 'accept:
application/hal+json' -d '{ "records": [
    { "svm": { "name": "svm1" }, "name" : "/vol/vol1/lun1", "space": {
"size": "1GB" }, "os_type": "linux" },
    { "svm": { "name": "svm1" }, "name" : "/vol/vol1/lun2", "space": {
"size": "5GB" }, "os_type": "linux" } ] }'
```

```
# The response
{
  "job": {
    "uuid": "09d0c372-a445-11eb-a707-005056bbbc41",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/09d0c372-a445-11eb-a707-005056bbbc41"
      },
      "results": {
        "href": "/api/storage/luns?job_results_uuid=09d0c372-a445-11eb-
a707-005056bbbc41"
      }
    }
  }
}
```

```
# GET the job results:
curl -X GET 'https://<mgmt-ip>/api/storage/luns?job_results_uuid=09d0c372-
a445-11eb-a707-005056bbbc41' -H 'accept: application/hal+json'
```

```
# The response:
{
  "records": [
    {
      "uuid": "2287871f-fbcf-4d64-ae89-59328dd755b3",
      "svm": {
        "name": "svm1"
      },
      "name": "/vol/vol1/lun1",
      "os_type": "linux",
      "space": {
        "size": 1073741824
      },
      "_links": {
        "self": {
          "href": "/api/storage/luns/2287871f-fbcf-4d64-ae89-59328dd755b3"
        }
      }
    }
  ],
  "num_records": 1,
  "errors": [
    {
      "message": "POST of record \"svm.name: svm1, name: /vol/vol1/lun2\"
failed. Reason: Volume offline.",
      "code": "262287"
    }
  ],
  "rollback_errors": [
    {
      "message": "DELETE of record \"uuid: 2287871f-fbcf-4d64-ae89-
59328dd755b3\" failed. Reason: Volume offline.",
      "code": "262287"
    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/luns?job_results_uuid=09d0c372-a445-11eb-a707-
005056bbbc41"
    }
  }
}
```

Record-based Errors

ONTAP Error Response Codes

| Error Code | Description | HTTP Code |
|---|---|---|
| 262287 | A single record in a bulk records request failed for the given reason. | 400 |
| 262288 | A record-based operation failed and was left in a partially completed state. | 400 |
| 262289 | A record-based job progress message. N of M records complete. | 200 |
| 262290 | A record-based job progress message during rollback. N of M records remaining. | 200 |
| 262291 | The operation did not complete within the system timeout period. Wait a few minutes, and then try the operation again. | 400 |
| 262292 | The records operation was partially completed. The following failures were encountered. | 400 |
| 262293 | Job is still running. Wait for the job to complete and then try the operation again. | 400 |
| 262294 | Job is not associated with this API. | 400 |

Option 2: Query-based PATCH and DELETE

APIs supporting PATCH or DELETE requests on a resource instance also support PATCH or DELETE on the collection, as long as at least one field is specified in the query portion of the URL. A PATCH or DELETE request issued on a collection is equivalent to internally doing a query-based GET, followed by a serial PATCH or DELETE operation on each matching record. However, it only does the operation for `return_timeout` seconds, which is 15 seconds, by default. If a query-based operation is not completed before `return_timeout` seconds, the API returns a next link. The client must use the next link with the same HTTP method to continue the operation. Query-based operations will not continue to the next record until the operation on the prior record is completed, even for operations that are normally asynchronous.

Example

```
# Modify the state of all volumes named "simpson" to be offline
PATCH /api/storage/volumes?name=simpson
{ "state": "offline" }
```

**Record filtering**

Overview

Records may either be filtered by performing queries that only apply to a single field at a time (though multiple of such queries may be done simultaneously for different fields), or by applying queries that search across a set of fields for a value fulfilling a single specified query.

Filtering records with single field queries

You can filter the results of a GET call using any attribute. The supplied query can either be for an exact value or can leverage special query operators.

```
<field>=<query value>
```

Filtering allows you to select objects where the specified field matches the supplied query, or which can contain wildcards, ranges, negations, or an OR-defined list of the above. The special query operators include the following:

Wildcard: `*`

```
abc*
abc*xyz

***xyz
```

Comparison: `< > <= >=`

```
<10
>=joe
```

Range: `..`

```
3..10
jim..joe
```

Negation: `!`

```
!3
!joe
!abc*
!jim..joe
```

Any of a list: `|`

```
3|5
3|5..9|>100
```

Escaping: `{} and ""`

The special query characters above can be treated literally, with no special meaning, by enclosing the value in either double quotes or curly brackets.
In order to escape double quotes or curly brackets and treat it literally you must enclose them with curly brackets.

```
"joe*"  # The double quotes escape the '*' and treat it as a literal
{a|b}   # The curly brackets escape the | and treat it as a literal
{{jim}} # The double curly brackets escape the curly brackets and treats
it as '{jim}'
{"abc"} # The curly brackets escape the double quotes and treats it as
'"abc"'
```

Filtering records with cross-field queries

Cross-field queries are useful when multiple fields should be searched for a value or some combination of values. Whereas traditional queries only allow a single field to be searched for a value, cross-field queries will return rows where any field in a specified set of fields matches the query. Cross-field queries may only be used for GET requests.

The fields to be queried across are specified in the "query_fields=" parameter. This should be a comma-delimited list of fields, or simply `*` to search across all fields.
To specify the query to use in the search, pass in the "query=" parameter to a GET request with the string to use as the query. Fields may also be excluded from searching prefixing with '!'. This is useful if all fields are specified with '*', and then certain fields wish to be excluded, or if an an entire object was queried, to exclude certain sub-fields.

Structure of the query

The query string represents a pattern to search for in all fields specified.

The `*` character is used to indicate wildcard character matching. `*` matches 0 or more of any character. For a query of "foo*bar", matches will include "foo123abcbar", "foobar", and "foo___123abcbar".
To search for any match among several possible patterns, the values may be ORed together with the '|' character. For example, to search for "foo" OR "bar", pass in "query=foo|bar". This may be extended to an arbitrary number of values, such as "query=foo|bar|baz".

Similarly, the query can be used to specify that multiple patterns must be found across all fields specified in "query_fields" for a row to be returned. To specify that multiple patterns must be found, include a space between each one. For example, to search across fields where the fields must contain both "foo" AND "bar", provide "query=foo bar". Again, this may be used on an arbitrary number of patterns. To search for rows that contain all of "foo" AND "bar" AND "baz" within the fields specified, provide "query=foo bar baz".

It should be noted that it is possible for all of the matches to a query to appear in a single field. For example, if "query=foo bar", and a field queried contains "foo bar blah", it will be considered a match. Obviously the queries matches can also be spread across different fields.

Examples

The following data is used for the examples below:

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 1 | widget1 | blue | chocolate | 1 2 3 | black cherry |
| 2 | widget2 | red | spinach | three fifty | maple |
| 3 | widget3 | rainbow | strawberry | thirty | spruce |
| 4 | widget4 | brown | strawberry chocolate | thirteen | willow |

**Request:** 'query_fields=color', 'query=red'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 2 | widget3 | red | spinach | three fifty | maple |

**Explanation:** The only row with a "color" column matching "red" is row 2.

**Request:** 'query_fields=id,number', 'query=3'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 1 | widget1 | blue | chocolate | 1 2 3 | black cherry |
| 3 | widget3 | rainbow | strawberry | forty two | spruce |

**Explanation:** Column "id" for row 3 matches the query, and column "number" for number for row 1 matches as well.

**Request:** 'query_fields=flavor', 'query=chocolate\|strawberry'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 1 | widget1 | blue | chocolate | 1 2 3 | black cherry |
| 3 | widget3 | rainbow | strawberry | forty two | spruce |
| 4 | widget4 | brown | strawberry chocolate | thirteen | willow |

**Explanation:** This query returns rows containing chocolate and/or strawberry in the flavor column. Rows 1, 2, and 4 all contain matches. Row 4 actually matches both queries.

**Request:** 'query_fields=flavor', 'query=chocolate strawberry'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 4 | widget4 | brown | strawberry chocolate | thirteen | willow |

**Explanation:** This query returns rows containing chocolate AND strawberry in the flavor column. Only row 4 contains matches for both queries.

**Request:** 'query_fields=name,number', 'query=*3\|three'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 1 | widget1 | blue | chocolate | 1 2 3 | black cherry |
| 2 | widget2 | red | spinach | three fifty | maple |
| 3 | widget3 | rainbow | strawberry | forty two | spruce |

**Explanation:** Searches across the name and number columns for a value either ending in '3', or containing "three". Row 1 contains 3 in the number field matching the first query, row 3 has a name of "widget3", matching the first query, and row 2 has a number containing three, matching the second query.

**Request:** 'query_fields=*', 'query=1\|2\|3 th*'

**Response:**

| id | name | color | flavor | number | tree |
|---|---|---|---|---|---|
| 2 | widget2 | red | spinach | three fifty | maple |
| 3 | widget3 | rainbow | strawberry | thirty | spruce |

**Explanation:** Searches across all columns, looking for rows where a row both contains either 1 and/or 2 and/or 3, and contains a value starting with "th". Row 1 contains a value matching 1 or 2 or 3, but has no column that begins with th. Similarly, row 4 has a value beginning with "th", but does not contain 1 or 2 or 3. Therefore only rows 2 and 3 are returned, which match both queries.

Cross-Field Query Errors

ONTAP Error Response Codes

| Error Code | Description | HTTP Code |
|---|---|---|
| 262272 | The specified query contains an unmatched quote. | 400 |
| 262273 | Both 'query_fields' and 'query' must be specified if either one is specified. | 400 |
| 262274 | The specified query is either empty or is equivalent to an empty query. | 400 |

| Error Code | Description | HTTP Code |
|---|---|---|
| 262274 | The parameters 'query_fields' and 'query' may only be specified for GET requests. | 400 |
| 262275 | At least one field must be specified for the cross-field query. | 400 |
| 262276 | A field was specified twice for the 'query_fields' parameter. | 400 |

**Requesting specific fields**

Overview

By default, calling GET on a collection generally returns only the properties that uniquely identify the record, along with a HAL 'self' link to the resource instance. However, you can choose the specific fields you want using the `fields` parameter. The `fields` parameter can also be used with GET when retrieving a single resource instance.

For discovery purposes, except for the CLI passthrough, the client can retrieve all standard properties using `fields=*`. These are the same properties returned when a GET is called on the specific instance using the path keys. However, using `fields=\*` is more expensive than selecting only the specific fields that are needed. In addition, because future releases may include additional properties in this list, or remove properties that were included by default, we strongly discourage using this in client-side software that is depending on specific fields being returned. To use the same list of fields when dealing with multiple versions of ONTAP, specify `ignore_unknown_fields=true`.

Some fields are more expensive to retrieve and are not included when using `fields=*` (or the instance-level GET). These fields are noted in the documentation. They can be returned either by specifying the fields directly, or by using `fields=\*\*`. However, we again strongly discourage this from being encoded into any client-side software. The performance of client software will suffer if a future version of ONTAP adds support for additional expensive properties.

```
fields=<field>[,...]
```

The `fields` input parameter allows you to specify exactly which fields you want to be returned.

Objects with fields

When an API contains fields that are objects, an entire object can be specified to return every field within the object. Individual fields within the object can be specified using dotted notation, as demonstrated below. Braces can be used to specify multiple fields within an object without repeating the object name. Braces can be nested within each other to select individual attributes within an object hierarchy.

> Dotted notation for arrays does not include array indices.

Examples

```
{
  "a": "<string>",
  "b": {
    "c": "<string>",
    "d": "<string>"
  },
  "e": [
    {
      "f": "<string1>",
      "g": "<string1>"
    },
    {
      "f": "<string2>",
      "g": "<string2>"
    }
  ]
}
```

Example `fields` query:

```
fields=a,b     // Fetch a, b.c, and b.d
fields=a,b.c,e // Fetch a, b.c, e.f, and e.g
fields=b.d     // Fetch b.d
fields=e.f     // Fetch e.f
fields=b,!b.c  // Fetch b.*, but not b.c
fields=b.{c,d} // Fetch b.c and b.d
```

**Records and pagination**

Records

Several query parameters control the return of records.

```
return_records=<true|false>
```

The default setting for `return_records` is `true` for GET calls and `false` for all other methods. When `false`, the array of records is not returned.

```
return_timeout=<0..120 seconds>
```

The `return_timeout` parameter specifies the number of seconds the cluster spends performing an operation before returning. The allowed range is `0` to `120` seconds. If the timeout is reached, GET calls return the records collected along with a pagination link. Other methods return and complete asynchronously. See

Non-blocking-operations for more details.

The default setting for `return_timeout` is `15` seconds for GET calls. For all other methods it is `0` seconds. This means that these calls might execute asynchronously in order to return as fast as possible.

> If the `order_by` parameter is specified, the operation might take longer because the collection is sorted before it is returned.

```
max_records=<number of records>
```

The `max_records` parameter limits the number of records that are returned (or acted on) before providing the "next" pagination link.

```
offset=<offset from the first record>
```

The `offset` parameter determines how many records to skip over prior to returning the first record.

For example, if you have a total of 15 records, and specify an offset of 10, only records 11-15 inclusive will be returned. When combined with a query or sorting specification, the offset will apply after the query or sorting, meaning that you will get records beginning at the Nth record, taking into account the query and sort order. Note that the computational cost of skipping over N records is likely as great as actually returning those N records.

Pagination

All calls to GET on a resource collection allow you to page through the results. If the `max_records` parameter is not specified, the cluster returns as many records as possible within the `return_timeout` time threshold. The number of records returned can be further limited by specifying a value for the `max_records` parameter. When the operation reaches either the `return_timeout` or the `max_records` threshold, it stops and returns the records as well as a HAL link that can be used to get the next page of records. It is possible for a pagination link to be returned even if there are no additional records. This occurs because the cluster does not check if there is an additional record before returning when it reaches a threshold. When there are potentially additional records, the response header will also contain a `Link` header containing the link followed by `rel="next"`.

The following is an example of the "next" link, which returns with a collection of records:

```
"_links": {
    "next": {
        "href":
"/api/storage/aggregates?start.aggregate=aggr25&max_records=25"
    }
}
```

Count only

The response to collection operations includes a `num_records` field. By passing `return_records=false`

with a GET call, you can retrieve the number of records without returning the records themselves. However, if either the `return_timeout` or `max_records` threshold is reached, an incomplete or partial number of records is returned and the "next" link must be called to retrieve additional record counts. All the partial counts must be added together to calculate the total count.

**Record sorting**

Overview

By default, records in a collection are returned in the order defined by the object. You can change the order by specifying the `order_by` query parameter. Most uses of the `order_by` parameter collect and reorder *all records* in the collection. This can be expensive when the collection is large. Therefore, clients are discouraged from paginating through the results with `max_records` when using `order_by`.

```
order_by=<field [asc|desc]>[,...]
```

If you want to sort on multiple fields (where the prior key value is the same), separate any fields (and optional direction) with a comma.

By default, sorting is done in ascending order based on the field type's ordering. If `desc` is specified after a field name, that field is sorted in descending order. Combining this with `max_records` allows you to see the `top` or `bottom` records based on the value of the specified field(s). When using this `top` or `bottom` functionality, queries on certain fields might require more time to search the entire collection regardless of the number of records actually returned.

Important Notes:

- When you use the `order_by` parameter, the `return_timeout` might be exceeded because the collection is sorted before it is returned.
- Using `order_by` on either a property of type array, or a nested property within an array (not including the records array), returns the records in an unspecified order.
- If the order_by value includes the direction (asc or desc), then you must make sure that the space between the field name and the direction is properly URL encoded when it is sent to the server. You may use either a `%20` or a `+` to encode the space, but if you send a space character, the server will respond with an error (400 Bad Request). Most programming language libraries will automatically do this encoding for you. Check the documentation of your language and/or library.

Examples

```
# Sort the volume collection from largest to smallest by size:
GET /api/storage/volumes?order_by=size+desc

# Find the top 5 applications using the most IOPS:
GET
/api/application/applications?order_by=statistics.iops.total+desc,name+asc
&max_records=5

# Find the top 10 applications using the most space and then

# if multiple applications are using the same space, sort them by IOPS:
GET
/api/application/applications?order_by=statistics.space.used+desc,statisti
cs.iops.total+desc&max_records=10
```

**Response body**

Overview

Every API call returns a top-level JSON object. These JSON objects includes GET calls that contain an array of records. This nesting technique allows metadata about the resource or resource collection to be returned as well as each resource instance.

GET calls that return an array of records can contain the following top-level elements:

```
{
  "records": [ {}, ... ]
  "num_records": <N>
  "_links": {
    "self": {
      "href": ...
    },
    "next": {
      "href": ...
    }
  }
}
```

- `records` - The array of records.
- `num_records` - The number of records in the array.
- `_links` - Links to relevant APIs, possibly including:
  - `self` - A link to retrieve the same data again.
  - `next` - If there are potentially more records, a link to retrieve the next page of records.

Custom response bodies

Some APIs might include additional top-level elements. For example, some APIs may include a top-level `errors` array which can include errors if the array of records is incomplete (for reasons other than pagination). See the documentation for each API to check for custom top-level elements.

Error objects

When an error occurs, an error object is returned in the response body. The error code is an integer returned in a JSON string. The optional `target` element is returned when ONTAP determines the error is due to a specific input field that you've supplied.

```
"error": {
  "message": "<string>",
  "code": "<integer>",
  "target": "<string>"
}
```

The following error codes can be returned by any ONTAP REST endpoint. Refer to individual endpoint documentation for additional endpoint specific error codes.

ONTAP Error Response Codes

| Error Code | Description | HTTP Code |
|---|---|---|
| 1 | An entry with the same identifiers already exists. | 409 |
| 2 | A field has an invalid value, is missing, or an extra field was provided. | 400 |
| 3 | The operation is not supported. | 405 |
| 4 | An entry with the specified identifiers was not found. | 404 |
| 6 | Permission denied. | 403 |
| 7 | Resource limit exceeded. | 429 or 503 |
| 8 | Resource in use. | 409 or 503 |
| 65541 | RPC timed out. | 500 |
| 65552 | Generic RPC failure. | 500 |
| 65562 | Internal RPC error | 500 |
| 262145 | Application code returned an unexpected exception. | 500 |
| 262160 | There are too many requests already being processed. Retry after a short delay. | 429 |
| 262177 | Missing value. | 400 |

| Error Code | Description | HTTP Code |
|---|---|---|
| 262179 | Unexpected argument. Argument shown in error message body. | 400 |
| 262185 | Invalid value with value in the body of the error. | 400 |
| 262186 | A field is used in an invalid context with another field, as shown in error message body. | 400 |
| 262188 | A field was specified twice. Location of assignments shown in error message body. | 400 |
| 262190 | You must provide one or more values to apply your changes. | 400 |
| 262194 | The expression is missing a value. | 400 |
| 262196 | Field cannot be set in this operation. | 400 |
| 262197 | Invalid value provided for field. Value and field shown in error message body. | 400 |
| 262198 | A request body is not allowed on GET, HEAD, and DELETE. | 400 |
| 262199 | Invalid JSON with error location provided in body of the error. | 400 |
| 262200 | Invalid JSON range, with range provided in the body of the error. | 400 |
| 262201 | Invalid JSON due to unknown formatting issue. | 400 |
| 262202 | Field is considered secret and should not be provided in the URL. | 400 |
| 262203 | Field is not supported in the body of a PATCH request. | 400 |
| 262210 | Unable to retrieve all required records within the timeout. This "order_by" query is not supported under the current system load with the current number of records in the collection. | 500 |
| 262211 | POST request on a REST API does not support filtering on an attribute. Attributes must be in the request body. | 400 |
| 262212 | Request is missing required value. | 400 |

| Error Code | Description | HTTP Code |
|---|---|---|
| 262220 | Wildcard fields=* is not allowed for CLI-based REST APIs. | 400 |
| 262222 | Specified value is listed more than once for this field. | 400 |
| 262245 | Invalid value with reason provided in body of the error. | 400 |
| 262247 | Invalid value for a field, with value and field in body of the error. | 400 |
| 262248 | A value is missing assignment operator. | 400 |
| 262249 | Field name is not supported for GET requests because it is not a readable attribute. | 400 |
| 262250 | Field name cannot be queried in a GET request because it is not a readable attribute. | 400 |
| 262254 | Invalid JSON input, an array was expected. | 400 |
| 262255 | An array was found in the JSON when it was not expected. | 400 |
| 262268 | The field is not supported as an order_by= field. | 400 |
| 262277 | The query_fields and query parameters may only be specified for GET requests. | 400 |
| 262282 | Property was specified twice. | 400 |
| 262286 | Mismatching braces found in the fields= query. | 400 |
| 263149 | Resource tag limit was reached while creating or modifying the resource. | 400 |
| 6691623 | User is not authorized for this operation. | 400 |
| 8586226 | Specified UUID not found. | 400 |
| 39387137 | A provided URL is invalid. | 400 |

### Synchronous and asynchronous operations

Overview

POST, PATCH, or DELETE operations that can take more than `2` seconds are considered *asynchronous* operations. They are implemented as non-blocking operations. Any API call that is expected to return in less than `2` seconds is considered *synchronous*. Synchronous operations ignore the `return_timeout` parameter.

API response

If the `return_timeout` is less than the time it takes for an operation to complete, the server returns the code `202 Accepted` after waiting for the specified `return_timeout` seconds. The default `return_timeout` for non-blocking operations is `0` seconds, meaning the operation returns as fast as possible. However, the operation never returns the success code `200 OK`, but instead returns either an error or the code `202 Accepted`.

The Location header

When a POST operation that is creating a resource returns `201 Created` (synchronous) or `202 Accepted` (asynchronous), the response header includes the `Location` of the resource. For asynchronous operations, a GET call on this resource link may return code `404 Not Found` until the operation successfully completes. Use the returned job link instead of the Location link to determine when the asynchronous operation is complete. POST operations that return code `200 OK` do not populate the Location header.

Example

The following example shows how to send a POST request and retrieve the Location header in return. Here "readonly" privileges are added to "/api/cluster" for the "test" security role:

```
# The API
POST "/api/security/roles/{owner.uuid}/{name}/privileges"
```

- "/api/security/roles/" is the whole collection of roles in the cluster.
- {owner.uuid} and {name} are the identifiers for a particular role and combined with the path to the collection form the URI for that role.
- The /privileges part of this API specifies that we want to operate on the sub-collection of privileges for a particular role.
- The POST request to this path will return a Location header that contains the URI for the newly created privilege. That URI can be accessed with GET in future calls to view the state of the privilege entry.

```
# The call:
curl -i -X POST -d '{"access":"readonly", "path":"/api/cluster"}'
"https://<mgmt-ip>/api/security/roles/001c3044-d15c-11ec-8091-
005056bb49e4/test/privileges"

# The response:
HTTP/1.1 201 Created
Date: Mon, 06 Jun 2022 15:11:42 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/security/roles/001c3044-d15c-11ec-8091-
005056bb49e4/test/privileges/%2Fapi%2Fcluster
Content-Length: 3
Content-Type: application/hal+json
```

The Location Header provides the location of the resource that is newly created as a part of the POST request. %2Fapi%2Fcluster is the identifier for the new privilege in the example above.

```
Doing a GET on the Location URL provided:

#The call:
curl -i -X GET "https://<mgmt-ip>/api/security/roles/001c3044-d15c-11ec-
8091-005056bb49e4/test/privileges/%2Fapi%2Fcluster"

# The response:
HTTP/1.1 200 OK
Date: Tue, 07 Jun 2022 17:04:50 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Length: 284
Content-Type: application/hal+json
Vary: Accept-Encoding
{
  "owner": {
  "uuid": "001c3044-d15c-11ec-8091-005056bb49e4"
  },
  "name": "test",
  "path": "/api/cluster",
  "access": "readonly",
  "_links": {
    "self": {
      "href": "/api/security/roles/001c3044-d15c-11ec-8091-
005056bb49e4/test/privileges/%2Fapi%2Fcluster"
    }
  }
}
```

Tracking non-blocking operations

Non-blocking or asynchronous operations are executed using *jobs*. The response to a non-blocking operation includes information about the job performing the operation, including a HAL link to the job resource. The job record also includes `state` and `message` fields. The `message` field indicates the progress of the operation while the `state` field indicates `running`. When a job is successful, the `state` and `message` fields indicate `success`. If an operation fails for any reason, the job's `state` reports `error`, and the `message` describes the problem that the operation encountered.

For POST operations, when a job is successfully completed, you can use the link from the `Location` header of the original response to retrieve the resource.

See GET /cluster/jobs

### HTTP status codes

Overview

The following supported HTTP status codes are returned by ONTAP:

- 200 OK: Returned for success when not creating a new object
- 201 Created: Returned for success after the creation of an object
- 202 Accepted: Returned when a job has been successfully started, but the operation is not complete
- 400 Bad Request: Returned if the input could not be parsed
- 401 Unauthorized: Returned if user authentication failed
- 403 Forbidden: Returned for authorization (RBAC) errors
- 404 Not Found: Returned when the specified resource does not exist
- 405 Method Not Allowed: Returned when the specified resource does not support the method (for example, POST or DELETE calls)
- 409 Conflict: Returned when there is a conflict with a different object that must be created, modified, or deleted before this operation can succeed
- 429 Too Many Requests: Returned when the client has sent more requests than the server can handle. The client should try again later (defined by the Retry-After header)
- 500 Internal Error: Returned for most other internal error codes
- 502 Bad Gateway: Returned if the application is temporarily unreachable. Try again later
- 503 Service Unavailable: Returned if the server is temporarily overloaded. Try again later.

### HTTP methods

Overview

The ONTAP REST API supports the following HTTP methods:

- GET: Supported on all collections to retrieve the records
- POST: When supported, calls on a collection to create the supplied resource
- PATCH: When supported, calls on a specific resource to update the supplied properties
- DELETE: When supported, calls on a specific resource to delete the resource
- HEAD: Supported wherever GET is supported. It makes a GET call, but only returns the HTTP headers
- OPTIONS: Supported on every endpoint so that you can determine which HTTP methods are supported

### Size properties

Overview

Many objects contain properties related to various sizes. Examples can be found in the `aggregate` object, `volume` object, `lun` object and `nvme_namespace` object. These properties are documented as type `integer`.

Unless otherwise documented, all sizes are reported in GET in bytes.

Depending on the development language-specific code generation, the API typically also requires an integer

value in bytes for POST and PATCH input as well.

Where a string value is accepted, such as query parameters and ad-hoc curl requests, any of the following suffixes can be used to specify different units:

| Suffix | Definition |
|---|---|
| KB | kilobytes (1024 bytes, aka kibibytes) |
| MB | megabytes (KB x 1024, aka mebibytes) |
| GB | gigabytes (MB x 1024, aka gibibytes) |
| TB | terabytes (GB x 1024, aka tebibytes) |
| PB | petabytes (TB x 1024, aka pebibytes) |

**SVM tunneling**

Overview

SVM tunneling allows for the scoping of REST APIs to any SVM from the cluster admin SVM interface. The HTTP headers "X-Dot-SVM-Name" and/or "X-Dot-SVM-UUID" are an alternative to supplying svm.name and/or svm.uuid in the request query or body. This allows for setting a context for an HTTP connection and reusing it for multiple calls. The cluster management interface or node management interface can be used instead of the desired SVM's interface. When either SVM tunneling or the SVM interface are used, any svm.name and svm.uuid properties need not be supplied in the request body, even if they are documented as required in the API. If they are supplied and do not match the SVM specified by the tunneling headers, then a 400 error will be returned.

Examples

Creates a new volume on SVM "vs0":

```
curl  -H "X-Dot-SVM-Name:vs0" -X POST "https://<mgmt-
ip>/api/storage/volumes" -d
'{"name":"vol1","aggregates":[{"name":"aggr1"}]}'
{
  "job": {
    "uuid": "b271e19d-c5cb-11e9-b97d-005056ac2211",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b271e19d-c5cb-11e9-b97d-005056ac2211"
      }
    }
  }
}
```

Retrieves all volumes on SVM "vs0":

```
curl -H "X-Dot-SVM-Name:vs0" -X GET "https://<mgmt-
ip>/api/storage/volumes"
{
  "records":[
  {
    "uuid":"a61e474-929a-4c78-882a-b72986ccf276",
    "name":"root_vs0",
    "_links":{
      "self":{
        "href":"/api/storage/volumes/aa61e474-929a-4c78-882a-b72986ccf276"
      }
    }
  },
  {
    "uuid":"b26c64f5-c5cb-11e9-b97d-005056ac2211",
    "name":"vol1",
    "_links":{
      "self":{
        "href":"/api/storage/volumes/b26c64f5-c5cb-11e9-b97d-005056ac2211"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self":{
    "href":"/api/storage/volumes"
    }
  }
}
```

Deletes a volume on SVM "vs0" using the X-Dot-SVM-UUID header:

```
curl -H "X-Dot-SVM-UUID:85ebedff-c43e-11e9-bc27-005056ac2211" -X DELETE
"https://<mgmt-ip>/api/storage/volumes?name=vol1"
{
  "jobs":[
    {
      "uuid":"4acf3f58-c5d2-11e9-b97d-005056ac2211",
      "_links":{
        "self":{
          "href":"/api/cluster/jobs/4acf3f58-c5d2-11e9-b97d-005056ac2211"
        }
      }
    }
  ],
  "num_records": 1,
  "_links":{
    "self":{
      "href":"/api/storage/volumes?name=vol1"
    }
  }
}
```

Retrieves all IP interfaces on SVM "vs3":

```
curl -H "accept: application/json" -H "X-Dot-SVM-Name:vs3" -X GET
"https://<mgmt-ip>/api/network/ip/interfaces"
{
  "records":[
    {
      "uuid":"83aeeac9-c5d8-11e9-b97d-005056ac2211",
      "name": "vs3_data_1"
    },
    {
      "uuid":"9c612bc0-c5a5-11e9-b97d-005056ac2211",
      "name":"vs3_data"
    }
  ],
  "num_records": 2
}
```

### Tagging Resources for Tracking Purposes

Overview

It may be useful for tools that use REST to categorize resources using tags for the purposes of tracking or grouping resources for certain purposes. For example, an application might want to create a set of resources

for a certain project or group and set a tag on all of those resources that indicates which project or group they were created for. The API provides a standardized property in some APIs for doing so called `_tags`. This property begins with an underscore to indicate that it is a generic property that is used for the same purpose across many different APIs, much like the `_links` property is.

Examples

Retrieve a list of all volumes that were provisioned for the staging environment:

```
curl -X GET "https://<mgmt-
ip>/api/storage/volumes?_tags=environment:staging"
{
  "records": [
    {"uuid": "a3af94b9-f33e-4fcd-80f7-96b8bb78229d", "name": "vol1",
"_tags": ["environment:staging"]},
    {"uuid": "6cbfef49-d234-4d37-8150-87ee15983696", "name": "vol2",
"_tags": ["environment:staging", "team:accounting"]},
    ...
  ]
}
```

Tag a volume as belonging to the accounting group when the volume is created

```
curl -X POST -d '{"name": "acct_vol", "svm": {"name": "svm1"},
"aggregates": [{"name": "aggr1"}], "_tags": ["team:accounting"]}'
"https://<mgmt-ip>/api/storage/volumes"
```

Remove an existing tag while leaving the rest

```
# Method 1: use the path to the volume as path key to the resource-tags
endpoint
curl -X DELETE "https://<mgmt-ip>/api/resource-
tags/team:accounting/resources/%2Fapi%2Fstorage%2Fvolumes%6cbfef49-d234-
4d37-8150-87ee15983696"

# Method 2: patch the volume directly

## get the current tags
curl -X GET "https://<mgmt-ip>/api/storage/volumes/6cbfef49-d234-4d37-
8150-87ee15983696?fields=_tags
{"uuid": "6cbfef49-d234-4d37-8150-87ee15983696", "name": "vol2", "_tags":
["environment:staging", "team:accounting"]}

## put back only the desired tags
curl -X PATCH -d '{"_tags": ["team:accounting"]}' "https://<mgmt-
ip>/api/storage/volumes/6cbfef49-d234-4d37-8150-87ee15983696"
```

Remove all tags

```
curl -X PATCH -d '{"_tags": []}' "https://<mgmt-
ip>/api/storage/volumes/6cbfef49-d234-4d37-8150-87ee15983696"
```

**Using the private CLI passthrough with the ONTAP REST API**

REST API access to CLI commands

To help CLI and ONTAP users transition to the ONTAP REST API, ONTAP provides a private REST API endpoint that can be used to access any CLI command. Usage of this API call is recorded and returned in the AutoSupport data collection so that NetApp can identify usability and functionality improvements in the REST API for future releases. There is no per-API documentation for the REST API access for each CLI command. Unlike the documented REST APIs, the API paths and properties for the CLI passthrough correspond very closely to the CLI. There are several rules that govern all the differences between a CLI command and the REST API mirroring the CLI command.

Rules for path differences when accessing a CLI command through the REST API

The API paths mirror the CLI paths, except for the use of the "show", "create", "modify", and "delete" verbs. Instead of using these four CLI verbs in the REST API, the corresponding HTTP methods must be used (GET, POST, PATCH, and DELETE). The four CLI verbs are removed from the API path supporting a command. For any commands where the last verb is hyphenated and begins with one of these verbs (for example, "show-space" or "delete-all"), you must remove the verb and following hyphen from the path. Any space in a full command path becomes a forward slash in the REST API (for example, "system node" becomes "/api/private/cli/system/node"). For non-show CLI commands that use non-standard verbs, the POST method should be used on the full path with the final verb in the API path. For example, "volume rehost" becomes "POST /api/private/cli/volume/rehost" and "cluster add-node" becomes "POST /api/private/cli/cluster/add-node".

To know which HTTP methods are supported for an API call, both documented and CLI-based, clients can use the "OPTIONS" HTTP method. For example, using OPTIONS on "/api/private/cli/volume" returns 'OK' with the HTTP "Allow" header containing a list of the supported HTTP methods (for example, "Allow: GET, HEAD, OPTIONS, POST, DELETE, PATCH"). For feature-specific CLI verbs, you can use OPTIONS on the API path. For example, using OPTIONS on "/api/private/cli/volume/restrict" returns with the HTTP header "Allow: OPTIONS, POST". Some of the CLI "show" commands do not contain the standard verb. For example, calling OPTIONS on "/api/private/cli/cluster/add-node-status" returns "Allow: GET, HEAD, OPTIONS".

There are some commands in the CLI that will not work using REST APIs. This includes most show commands that do not support "show -fields" in the CLI. The REST API also does not support CLI commands that create a new shell (like "run" and "vserver context").

Here are several examples of mappings from the ONTAP CLI to the ONTAP REST API for the /api/private/cli path:

- volume show → GET /api/private/cli/volume
- volume create → POST /api/private/cli/volume
- volume modify → PATCH /api/private/cli/volume
- volume delete → DELETE /api/private/cli/volume
- volume restrict → POST /api/private/cli/volume/restrict
- volume show-space → GET /api/private/cli/volume/space
- volume show-footprint → GET /api/private/cli/volume/footprint
- cluster add-node → POST /api/private/cli/cluster/add-node
- cluster add-node-status → GET /api/private/cli/system/node/add-node-status
- system node coredump show → GET /api/private/cli/system/node/coredump
- system node coredump delete → DELETE /api/private/cli/system/node/coredump
- system node coredump delete-all → DELETE /api/private/cli/system/node/coredump/all

Rules for field differences when accessing a CLI command through the REST API

All CLI parameters are supported in the CLI-based REST APIs. However, REST converts hyphens (-) in CLI parameter names to underscores (_) in the REST API JSON response body. In general, REST API responses use the same formatting for property values as ONTAPI. For example, enumerated values are formatted in lowercase instead of uppercase and with underscores instead of hyphens in the REST API response body. Both CLI and ONTAPI formats are allowed on input. Also similar to ONTAPI, sizes and percentages in REST are encoded as integers in bytes. Unlike ONTAPI or the CLI, date and time values in REST are encoded with the ISO-8601 format. All fields that you want returned from the GET call must be specified using the `fields` parameter. Note that the /api/private/cli/… APIs do not support "fields=*".

Examples

Retrieve OPTIONS for volumes endpoint (with results contained in header):

```
curl -X OPTIONS "https://<mgmt-ip>/api/private/cli/volume" --include
Allow: GET, HEAD, OPTIONS, POST, DELETE, PATCH
{
}
```

GET size and percent-used for all volumes:

```
curl -X GET "https://<mgmt-ip>/api/private/cli/volume?fields=size,percent-
used&pretty=false"
{
  "records": [
    { "vserver": "vs1", "volume": "vol1", "size": 20971520,
"percent_used": 73 },
    { "vserver": "vs1", "volume": "vol2", "size": 20971520,
"percent_used": 87 },
    ...
  ]
}
```

GET size and percent-used for a specific volume:

```
curl -X GET "https://<mgmt-
ip>/api/private/cli/volume?volume=vol2&pretty=false"
{
  "records": [
    { "vserver": "vs1", "volume": "vol2", "size": 209715203864,
"percent_used": 89 },
    ...
  ]
}
```

POST a new volume with all required attributes:

```
curl -X POST "https://<mgmt-ip>/api/private/cli/volume" -d
'{"volume":"vol3","vserver":"vs0","aggregate":"aggr1"}'
{
  "job": {
    "uuid": "f7b5f5cb-54a2-11e9-930a-005056ac6a3f",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f7b5f5cb-54a2-11e9-930a-005056ac6a3f"
      }
    }
  },
  "cli_output": "[Job 36] Job is queued: Create vol2."
}
```

Attempt to DELETE an online volume:

```
curl -X DELETE "https://<mgmt-
ip>/api/private/cli/volume?vserver=vs1&volume=vol1"
{
  "num_records": 0,
  "error: {
    "message": "Volume vol1 in Vserver vs1 must be offline to be
deleted.",
    "code": "917658"
  }
}
```

PATCH a volume to become offline:

```
curl -X PATCH "https://<mgmt-
ip>/api/private/cli/volume?vserver=vs1&volume=vol1" -d '{ "state":
"offline" }'
{
  "num_records": 1,
  "cli_output: "Volume modify successful on volume vol1 of Vserver vs1.\n"
}
```

DELETE the offline volume:

```
curl -X DELETE "https://<mgmt-
ip>/api/private/cli/volume?vserver=vs1&volume=vol1"
{
  "jobs": [
    {
      "uuid": "3f35a934-4b40-11e9-9f4d-005056bbf4eb",
      "_links": {
        "self": {
          "href": "/api/cluster/jobs/3f35a934-4b40-11e9-9f4d-005056bbf4eb"
        }
      }
    }
  ],
  "num_records": 1,
  "cli_output": "[Job 1243] Job succeeded: Successful\n"
}
```

When POST is called for a command that uses a job, the REST API does not wait for the job to complete, unless `return_timeout` is specified. However, PATCH and DELETE calls on the command path (using queries on key fields in the query portion of the URI) wait up to 15 seconds for the operation to complete if the `return_timeout` parameter is not specified.

DELETE an offline volume without waiting:

```
curl -X DELETE "https://<mgmt-
ip>/api/private/cli/volume?vserver=vs1&volume=vol2&return_timeout=0"
{
  "jobs": [
    {
      "uuid": "a7138c5e-4b69-11e9-9f4d-005056bbf4eb",
      "_links": {
        "self": {
          "href": "/api/cluster/jobs/a7138c5e-4b69-11e9-9f4d-
005056bbf4eb",
        }
      }
    }
  ],
  "num_records": 1,
  "cli_output": "[Job 1247] Job is queued: Delete vol1.\n"
}
```

CLI message output

As shown in the previous example, any non-field and non-error based output that would have appeared in the CLI is returned in a top-level `cli_output` attribute in the response body. This does not contain normal CLI headers or field values. It only displays messages that were printed to the CLI.

HTTP status codes

Error codes in the response body are mapped to the most appropriate HTTP status codes. In cases where this is not done, the HTTP status code defaults to 500. This does not necessarily indicate that the error is internal to ONTAP.

Security

All CLI-based REST APIs are RBAC-controlled, based on the role of the authenticated user and have the same protections they have in the CLI.

Location of CLI fields for CLI-based REST APIs:

- POST APIs: All CLI fields must be provided in the request body.
- GET APIs: All desired CLI fields (except keys) must be specified in the `fields` parameter. The non-key fields returned via the CLI will not be returned if not requested. The client can also provide a query for any field.

- PATCH APIs: The client can provide a query for any field, but at least one field must have a query. To modify only a single record, all CLI keys must contain an exact query. All new values for the object must be provided in the request body.
- DELETE APIs: The client can provide a query for any field, but at least one field must have a query. To delete only a single record, all CLI keys must contain an exact query. Non-attribute inputs (such as `force`) must be provided in the query portion of the URI.

# Application

## ONTAP REST API Application endpoints

### Overview

ONTAP application APIs simplify storage management by using terminology specific to a type of application. This application-specific terminology can be used to provision and manage ONTAP storage objects. A single call using application-specific parameters provisions storage and enables protocol access for an application following NetApp best practices. You can view and manage the ONTAP objects making up the application as a group using the application APIs. The library of available application templates already includes several database and virtualization applications.

### APIs

There are several application APIs that must be used to fully manage an application. Templates are used to represent any parameters specific to a given application. Some APIs expose applications in terms of their specific template, while others only expose a generic view that all applications share. The template view is present on the *templates* and *applications* APIs (although these APIs do also include some generic fields). The *components* and *snapshots* APIs are entirely generic and do not differ across types of applications.

The following section provides an overview of each API, followed by a lifecycle example of managing an application to demonstrate how the APIs can be used together.

---

#### Template

A template is an ONTAP representation of a specific type of application. Each template represents one type of application, the parameters that can be used to customize it, the layout of its storage, and how it can be accessed. Templates are intended to expose an application in terms specifically applicable to an administrator of a given application. As such, traditional ONTAP storage elements are generally not included in an application template.

The template APIs can be used to discover what templates are currently available. The ONTAP API documentation also includes a model of the templates. The template APIs generally provide the same information as the documentation, but the template APIs might provide more up-to-date details about the default values of template parameters based on the current ONTAP configuration. However, only the ONTAP API documentation includes a full description of each template parameter, its usage, and whether it is optional.

---

#### Application

The application APIs are the only interfaces that allow management of an application using template properties.

---

The application object includes the following three sections:

1. Generic metadata about the application, including common fields such as the name of the application, the template used to provision it, and the generation number of the application.
2. Statistics information about the application, including space and IOPS details about the entire application and each of its components. These are expensive to collect and should only be requested when needed using a *fields*= query.
3. A template view of the application. The application object itself presents a mutually exclusive list of all possible templates. Only one of these fields can be used per application. The name of the field corresponds to the name of the template used by the application. Currently, the creation of a new application and the modification of the storage service for an existing application are supported through the template parameters.

---

#### Component

The component API offers a generic view of the application and how to access the application from the host application. This is the only API that exposes the underlying ONTAP storage elements of which the application is composed. It is read-only; it cannot support modifications specific to the type of application it is presenting.

The component object includes the following details for an application:

1. The NFS export rules for accessing the application from the host.
2. The CIFS share and users that can access the application from the host.
3. For IP-based protocols, the IP addresses that are best suited for accessing the component.
4. The underlying storage elements that make up the component, such as volumes or LUNs.
5. The SAN initiators that can access the application from the host.

---

#### Snapshot

The snapshot APIs offer full CRUD for application-level snapshots. Application snapshots can be flagged as either crash-consistent or application-consistent. From the perspective of ONTAP, there is no difference between the two. It is the responsibility of the administrator to ensure that the application is in a consistent state before flagging a snapshot as application-consistent. Use of the SnapCenter Backup Management suite is recommended to ensure correct interaction between host applications and ONTAP.

---

### Example

The following example outlines the APIs necessary to manage applications and how they fit together. However, this example does not provide detailed information on each API. See the documentation for the individual APIs for more information.

#### 1) Discover the templates

This documentation, which includes the model of each template as part of the *templates* and *applications* APIs, is the easiest and most comprehensive way to discover the available templates. The *templates* API can also be used to query the system for templates in a programmatic way.

To discover the templates available to provision an Oracle application, the following query is used.

```
# The API:
/api/application/templates

# The query:
name=oracle*

# The call:
curl -X GET "https://<mgmt-ip>/api/application/templates?name=oracle*" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "name": "oracle_on_nfs", "description": "Oracle using NFS."
  },

  {
    "name": "oracle_on_san", "description": "Oracle using SAN."
  },

  {
    "name": "oracle_rac_on_nfs", "description": "Oracle RAC using NFS."
  },
  {
    "name": "oracle_rac_on_san", "description": "Oracle RAC using SAN."
  }

],
"num_records": 4
}
```

#### 2) Create an application

Now that we know the possible templates, we use one to create an application. The template properties differ from template to template, and can be found by exploring the model of the application object in this documentation. Each call to create an application must include the properties for exactly one template. These properties are provided under the property with the same name as the template. Other than the template properties, the only other required properties to create an application are the SVM and name.

In the following call example, not all of the template properties are included. Where a property is not needed or the default is sufficient, the property can be excluded. In this case using the *oracle_on_nfs* template, the *archive_log*, and *protection_type* are not included. The template name, *oracle_on_nfs*, is specified above the group of template properties, after the names of the application and the SVM.

Creating an application is asynchronous, so the response for this API includes information about the job doing the work. The response header also includes the *location* of where the application can be found if the job is successful.

---

Prior to creating an application, the following prerequisites must be met for the protocols associated with the template:

- Licences must be installed.
  - POST /cluster/licensing/licenses
- Aggregates must exist with enough available space and IOPS to satisfy the requested size.
  - POST /storage/aggregates
- An SVM must exist with protocol services enabled.
  - POST /svm/svms
- LIFs must exist. For SAN applications, only High Availability groups where each node has at least one LIF will be considered for placement of storage objects.
  - POST /network/ip/interfaces
  - POST /network/fc/interfaces

---

The following are not required prior to creating an application, but might be necessary before connecting to the application:

- Network routes must be created to access ethernet based LIFs.
  - POST /network/ip/routes
- For volumes created by this operation to be successfully mounted, ONTAP requirements related to mounting must be met.

---

```
# The API:
/api/application/applications

# The query:
No query is needed for this command. Optionally, you can specify the
return_timeout or set the return_records flag to alter the behavior of the
command.

# The body:
```

```
{
"name": "my_ora_app",
"svm": {
  "name": "svm1"
},
"oracle_on_nfs": {
  "db": {
    "size": "2GB",
    "storage_service": {
      "name": "value"
    },
    "nfs_access": [
        {
            "access": "rw",
            "host": "0.0.0.0/0"
        }
    ]
  },
  "redo_log": {
    "size": "1GB"
  },
  "ora_home": {
    "size": "1GB"
  }
}
}

# The call:
    curl -X POST "https://<mgmt-ip>/api/application/applications" -H
"accept: application/hal+json" -H "content-type: application/json" -d '{
"name": "my_ora_app", "svm": { "name": "vs1" }, "oracle_on_nfs": { "db": {
"size": "2GB", "storage_service": { "name": "value" }, "nfs_access": [ {
"access": "rw", "host": "0.0.0.0/0" } ] }, "redo_log": { "size": "1GB" },
"ora_home": { "size": "1GB" } } }'

# The response:
{
"job": {
  "uuid": "dc0d01dd-df5a-11e7-b5d2-005056b47eb2",
  "id": 94,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/dc0d01dd-df5a-11e7-b5d2-005056b47eb2"
    }
  }
}
```

```
}

# The response header:
date: Tue, 12 Dec 2017 16:38:18 GMT
server: libzapid-httpd
content-type: application/hal+json
location: /api/application/applications/dbc10d87-df5a-11e7-b5d2-
005056b47eb2
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
keep-alive: timeout=5, max=100
content-length: 203
```

#### 3) Wait for the application to be created

The call to create the application returns information about the job, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete, and the application has been fully created, the message indicates success and the *state* of the job property is *success*.

For brevity purposes, the successful job response is shown here. On a real cluster, an application might take several seconds to several minutes to be created, depending on the system load. If the job is not complete, the *message* property includes a short description on the progress of the job, and the *state* indicates *running*.

```
# The API:
/api/cluster/jobs/{uuid}

# The call, provided by the HAL link from step 3:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/dc0d01dd-df5a-11e7-b5d2-
005056b47eb2" -H "accept: application/hal+json"

# The response:
{
"uuid": "dc0d01dd-df5a-11e7-b5d2-005056b47eb2",
"state": "success",
"message": "Complete: Success [0]",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/dc0d01dd-df5a-11e7-b5d2-005056b47eb2"
  }
}
}
```

#### 4) Retrieve the new application

You can look up the application directly without listing all the applications. Use the *location* header that is included in the response when the application is created.

> The following example uses a query to retrieve only a small number of the application's properties.

```
# The API:
/api/application/applications/{uuid}

# The query:
fields=name,template.name,generation,state

# The call:
curl -X GET "https://<mgmt-ip>/api/application/applications/dbc10d87-df5a-
11e7-b5d2-005056b47eb2?fields=name,template.name,generation,state" -H
"accept: application/json"

# The response:
{
"uuid": "dbc10d87-df5a-11e7-b5d2-005056b47eb2",
"name": "my_ora_app",
"template": { "name": "oracle_on_nfs" },
"generation": 2,
"state": "online"
}
```

#### 5) Discover how to access the application

The components API provides information on how to access the storage that is provisioned for the application.

For brevity, only the names of the components are requested. See the API documentation for more information on the other available fields.

```
# The API:
api/application/applications/{application.uuid}/components

# The query:
fields=name

# The call:
curl -X GET "https://<mgmt-ip>/api/application/applications/dbc10d87-df5a-
11e7-b5d2-005056b47eb2/components?fields=name" -H "accept:
application/json"

# The response:
{
"records": [
  { "uuid": "e06fb407-df5a-11e7-b5d2-005056b47eb2", "name": "db" },
  { "uuid": "e0709732-df5a-11e7-b5d2-005056b47eb2", "name": "ora_home" },
  { "uuid": "e07158eb-df5a-11e7-b5d2-005056b47eb2", "name": "redo_log" }
],
"num_records": 3
}
```

---

#### 6) Update the application

To update the storage service, the same template that is used for creating the application is reused, but with only the storage_service properties set. In the generic SAN and NAS templates, the name of each component must also be specified.

In this example, the cluster only supports the *value* storage service, so modifications of the application to a faster storage service fail. Note how the error message indicates the parameter that caused the problem.

Application modification, like application creation, is an asynchronous operation. If a valid command is passed, the API returns information about the job instead of an error.

---

```
# The API:
/api/application/applications/{uuid}

# The body:
{
"oracle_on_nfs": { "db": { "storage_service": { "name": "extreme" } } }
}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/application/applications/dbc10d87-
df5a-11e7-b5d2-005056b47eb2" -H "accept: application/hal+json" -H
"content-type: application/json" -d '{ "oracle_on_nfs": { "db": {
"storage_service": { "name": "extreme" } } } }'

# The response:
{
"error": {
  "message": "Invalid value for parameter \"oracle_on_nfs.db.storage-
service.name\": extreme. Supported values are: value.",
  "code": "65995152"
}
}
```

---

#### 7) Manage snapshots

For applications created with the *local protection_type* set to *hourly*, snapshots are automatically taken every hour. These snapshots can be retrieved or restored using the snapshot APIs. Snapshots can also be taken on demand using these APIs. It is important to note that the *consistency_type* flag of the snapshot is for record-keeping only: it is the responsibility of the administrator to ensure that the application is in a consistent state prior to flagging a snapshot as *application* consistent.

Take a snapshot manually:

---

```
# The API:
/api/application/applications/{uuid}/snapshots

# The body:
{
  "name": "little_bobby_tables",
  "consistency_type": "crash"
}

# The call:
curl -X POST "https://<mgmt-ip>/api/application/applications/dbc10d87-
df5a-11e7-b5d2-005056b47eb2/snapshots" -H "accept: application/hal+json"
-H "content-type: application/json" -d '{ "name": "little_bobby_tables",
"consistency_type": "crash"}'

# The response:
{}

# The response header:
date: Tue, 12 Dec 2017 17:40:10 GMT
server: libzapid-httpd
content-type: application/hal+json
location: /api/application/applications/dbc10d87-df5a-11e7-b5d2-
005056b47eb2/snapshots/dbc10d87-df5a-11e7-b5d2-
005056b47eb2_13_little_bobby_tables
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
keep-alive: timeout=5, max=100
content-length: 3
```

In the above example, the response body is empty, and the response header includes the *location* of the newly created snapshot. By default, all POST calls return an empty body unless a job is used to process the creation asynchronously. This behavior can be changed with the query flag *return_records*.

Restoring a snapshot uses an action API. Action paths can also be performed asynchronously as jobs, as with creating or modifying an application. The response header does not include a *location*, because this action is not creating a resource.

```
# The API:
/api/application/applications/{application.uuid}/snapshots/{snapshot.uuid}
/restore

# The call:
curl -X POST "https://<mgmt-ip>/api/application/applications/dbc10d87-
df5a-11e7-b5d2-005056b47eb2/snapshots/dbc10d87-df5a-11e7-b5d2-
005056b47eb2_13_little_bobby_tables/restore" -H "accept:
application/hal+json"

# The response:
{
"job": {
  "uuid": "00e81690-df64-11e7-b5d2-005056b47eb2",
  "id": 100,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/00e81690-df64-11e7-b5d2-005056b47eb2"
    }
  }
}
}

# The response header:
date: Tue, 12 Dec 2017 17:43:46 GMT
cache-control: no-cache,no-store,must-revalidate
server: libzapid-httpd
connection: Keep-Alive
keep-alive: timeout=5, max=100
content-length: 204
content-type: application/hal+json
```

---

### Smart containers

Smart containers are traditional ONTAP storage objects such as FlexVol or FlexGroup created using the application REST API.

- NAS - FlexVolume, FlexGroup, and FlexCache objects
- SAN - LUNs
- NVME - Namespaces
- S3 - Object Store S3 Buckets

---

The benefits of creating a Smart Container are as follows:

- ONTAP determines the best placement for the storage object based on available performance and space capacity.
- Access controls can be optionally set.
- Snapshot schedules can be optionally set.
- A single atomic job that does all the above.

---

Smart containers are similar to generic enterprise applications (NAS, SAN, NVME), but with certain restrictions. Smart containers are restricted to one application-component. Any post-provisioning data management operations on smart containers must be performed via PATCH operations corresponding to the object created. However, the POST, GET and DELETE operations that exist for applications will also operate for smart containers.

To create a Smart Container the "smart_container:true" parameter must be provided.

---

Prior to creating a smart container, the following prerequisites must be met for the protocols associated with the template:

- Licences must be installed.
  - POST /cluster/licensing/licenses
- Aggregates must exist with enough available space to satisfy the requested size.
  - POST /storage/aggregates
- An SVM must exist with protocol services enabled.
  - POST /svm/svms
- LIFs must exist. For SAN objects, only High Availability groups where each node has at least one LIF to be considered for placement of storage objects.
  - POST /network/ip/interfaces
  - POST /network/fc/interfaces

---

The following are not required prior to creating a smart container:

- Network routes must be created to access Ethernet-based LIFs.
  - POST /network/ip/routes
- To mount volumes by this operation successfully, all ONTAP requirements related to mounting must be met.

---

**Example**

The following examples outline the APIs necessary to create a smart container. Two types of smart container creation are supported:

- A smart container with new ONTAP storage objects as specified in the JSON body.
- An existing ONTAP volume can be converted into a smart container (supported only on generic SAN and NVME templates).

This is an addendum to the example provided on how to create an application.

```
# The API:
/api/application/applications

# The query:
No query is needed for this command. Optionally, you can specify the
return_timeout or set the return_records flag to alter the behavior of the
command.

# The body:
Creates a smart container with new ONTAP storage objects:
{
"name": "my_container",
"svm": {
    "name":"vs1"
},
"template": {
    "name":"nas"
},
"smart_container": "true"
"nas": {
    "application_components": [
        {
            "share_count": "1",
            "name": "myVolume",
            "storage_service": {
                "name": "value"
            },
            "total_size": "100mb"
        }
    ]
}
}

Converting an existing volume into a smart container:
{
"name": "my_container",
"svm": {
    "name":"vs1"
},
```

```
"template": {
    "name":"san"
},
"smart_container": "true"
"san": {
    "application_components": [
        {
            "name": "existingVolume" #name of an existing volume
        }
    ]
}
}


# The call:
Creates a smart container with new ONTAP storage objects:
curl -X POST "https://<mgmt-ip>/api/application/applications" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d '{"name":
"my_container", "svm": {"name": "vs1"} , "smart_container": true ,
"template": {"name": "nas"} , "nas": {"application_components":
[{"share_count": "1", "name": "myVolume", "storage_service": {"name":
"value"} , "total_size": "100mb"} ] } }'

Converting an existing volume into a smart container:
curl -X POST "https://<mgmt-ip>/api/application/applications" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d '{"name":
"my_container", "svm": {"name": "vs1"} , "smart_container": true ,
"template": {"name": "san"} , "san": {"application_components": [{"name":
"existingVolume"} ] } }'

# The response:
 {
   "job": {
       "uuid": "5440db05-77f0-11e9-a5a0-005056bba32f",
       "_links": {
           "self": {
               "href": "/api/cluster/jobs/5440db05-77f0-11e9-a5a0-
005056bba32f"
           }
       }
   }
 }


# The response header:
date: Tue, 23 May 2019 16:38:18 GMT
```

```
server: libzapid-httpd
content-type: application/hal+json
location: /api/application/applications/5440db05-77f0-11e9-a5a0-
005056bba32f
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
keep-alive: timeout=5, max=100
content-length: 203
```

#### Smart Container Properties

- exclude_aggregates - This property is available for SAN, NAS and S3 Smart Containers. The list of aggregates specified in this property will be excluded while provisioning the storage object. If no suitable aggregate can be found to place the storage object, smart container provisioning will fail.

These sections are only allowed for smart containers and will return an error when provided on traditional applications. The following is an example of the error returned:

```
{
"error": {
  "message": "Field \"<field>\" is only supported on smart containers.",
  "code": "65996161"
  }
}
```

#### Updating the smart container

A smart container can be updated to add more LUNS and/or namespaces, with the same template used to create a smart container being reused, and with the following fields set:

- lun_count/namespace_count - represents the total number of LUNS/namespaces in the smart container.
- os_type - represents the OS type of the new LUNS/namespaces.
- total_size - represents the total size of the new LUNS/namespaces to be added.
- igroup_name/subsystem - represents the igroup/subsystem mapping for the new LUNS/namespaces. Updates are allowed only on generic SAN and NVME templates.

```
# The API:
/api/application/applications/{uuid}

# The body:
{
"san": {
    "application_components": [
        {
            "name": "myVolume",
            "lun_count": 4
            "total_size": "1gb"
            "os-type": "linux",
            "igroup_name": "igroup1"
        }
    ]
}
}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/application/applications/dbc10d87-
df5a-11e7-b5d2-005056b47eb2" -H "accept: application/hal+json" -H
"content-type: application/json" -d '{ "san": { "application_components":
[{ "name": "myVolume", "total_size": "1GB", "lun_count": 4, "os_type":
"linux", "igroup_name": "igroup1"}]}}'
```

## Application API limitations

### Template versus generic

Applications can be represented in either template or generic terms. All applications can be represented in generic terms as a list of components. Each component generally maps to a field in the template. For example, Microsoft SQL Server applications have a component named *sqldata* that corresponds to the *db* parameter in the sql_on_san template. These mappings are usually straightforward and allow the templates to present application terminology, while the generic view uses the traditional naming schemes for ONTAP storage elements.

The current release supports the creation and modification of applications in template terms, but retrieval is not supported. The mapping from template to generic terms is left to your own discretion.za

### ONTAP feature support

Application APIs are interfaces layered on top of traditional ONTAP storage. While the intent is to provide a full management suite through application APIs, some features of the underlying ONTAP objects are not directly supported through application APIs. Applications are provisioned using ONTAP best practices, so the need for additional modifications of the underlying objects should be minimal. If such modifications are necessary, the traditional ONTAP APIs can be used. The */api/application/{application.uuid}/components* API provides a *backing_storage* field that can be used to locate the storage objects associated with an application. This API also provides details of the NFS, SAN, or CIFS protocol access objects associated with the application.

The application APIs use the extra information known about the application to coordinate multiple ONTAP objects in unison. When using non-application APIs, certain settings might interfere with the ONTAP object coordination and cause the application APIs to behave unexpectedly. To continue to supply the full ONTAP feature set, these modifications on the underlying objects are allowed, but there is no guarantee that these modifications will not adversely affect the application experience. You should use this feature with caution.

## Retrieve applications

GET `/application/applications`

**Introduced In:** 9.6

Retrieves applications.

### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `<template>` the property corresponding to the `template.name` of the application

### Query examples

Numerous queries are available for classifying and sorting applications:

1. Return a list of applications sorted by name.

```
GET /application/applications?order_by=name
```

2. Return a list of applications for a specific SVM.

```
GET /application/applications?svm.name=<name>
```

3. Return a list of all SQL applications.

```
GET /application/applications?template.name=sql*
```

4. Return a list of all applications that can be accessed via SAN.

```
GET /application/applications?template.protocol=san
```

1. Return the top five applications consuming the most IOPS.

```
  GET /application/applications?order_by=statistics.iops.total
desc&max_records=5
```

The above examples are not comprehensive. There are many more properties available for queries. Also, multiple queries can be mixed and matched with other query parameters for a large variety of requests. See the per-property documentation below for the full list of supported query parameters.

**Learn more**

- DOC /application

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | query | False | Filter by UUID |
| name | string | query | False | Filter by name |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| template.name | string | query | False | Filter by template.name |
| template.version | string | query | False | Filter by template.version |
| template.protocol | string | query | False | Filter by template.protocol |
| generation | string | query | False | Filter by generation |
| state | string | query | False | Filter by state |
| protection_granularity | string | query | False | Filter by protection granularity |
| rpo.is_supported | string | query | False | Filter by rpo.is_supported |
| rpo.local.name | string | query | False | Filter by rpo.local.name |
| rpo.local.description | string | query | False | Filter by rpo.local.description |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rpo.remote.name | string | query | False | Filter by rpo.remote.name |
| rpo.remote.description | string | query | False | Filter by rpo.remote.description |
| rpo.components.name | string | query | False | Filter by rpo.components.name |
| rpo.components.uuid | string | query | False | Filter by rpo.components.uuid |
| rpo.components.rpo.local.name | string | query | False | Filter by rpo.components.rpo.local.name<br>• Introduced in: 9.7 |
| rpo.components.rpo.local.description | string | query | False | Filter by rpo.components.rpo.local.description |
| rpo.components.rpo.remote.name | string | query | False | Filter by rpo.components.rpo.remote.name |
| rpo.components.rpo.remote.description | string | query | False | Filter by rpo.components.rpo.remote.description |
| statistics.space.provisioned | string | query | False | Filter by statistics.space.provisioned |
| statistics.space.used | string | query | False | Filter by statistics.space.used |
| statistics.space.used_percent | string | query | False | Filter by statistics.space.used_percent |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.space.used_excluding_reserves | string | query | False | Filter by statistics.space.used_excluding_reserves |
| statistics.space.logical_used | string | query | False | Filter by statistics.space.logical_used |
| statistics.space.reserved_unused | string | query | False | Filter by statistics.space.reserved_unused |
| statistics.space.available | string | query | False | Filter by statistics.space.available |
| statistics.space.savings | string | query | False | Filter by statistics.space.savings |
| statistics.iops.total | string | query | False | Filter by statistics.iops.total |
| statistics.iops.per_tb | string | query | False | Filter by statistics.iops.per_tb |
| statistics.snapshot.reserve | string | query | False | Filter by statistics.snapshot.reserve |
| statistics.snapshot.used | string | query | False | Filter by statistics.snapshot.used |
| statistics.latency.raw | string | query | False | Filter by statistics.latency.raw |
| statistics.latency.average | string | query | False | Filter by statistics.latency.average |
| statistics.statistics_incomplete | string | query | False | Filter by statistics.statistics_incomplete |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.shared_storage_pool | string | query | False | Filter by statistics.shared_storage_pool |
| statistics.components.name | string | query | False | Filter by statistics.components.name |
| statistics.components.uuid | string | query | False | Filter by statistics.components.uuid |
| statistics.components.storage_service.name | string | query | False | Filter by statistics.components.storage_service.name |
| statistics.components.space.provisioned | string | query | False | Filter by statistics.components.space.provisioned |
| statistics.components.space.used | string | query | False | Filter by statistics.components.space.used |
| statistics.components.space.used_percent | string | query | False | Filter by statistics.components.space.used_percent |
| statistics.components.space.used_excluding_reserves | string | query | False | Filter by statistics.components.space.used_excluding_reserves |
| statistics.components.space.logical_used | string | query | False | Filter by statistics.components.space.logical_used |
| statistics.components.space.reserved_unused | string | query | False | Filter by statistics.components.space.reserved_unused |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.components.space.available | string | query | False | Filter by statistics.components.space.available |
| statistics.components.space.savings | string | query | False | Filter by statistics.components.space.savings |
| statistics.components.iops.total | string | query | False | Filter by statistics.components.iops.total |
| statistics.components.iops.per_tb | string | query | False | Filter by statistics.components.iops.per_tb |
| statistics.components.snapshot.reserve | string | query | False | Filter by statistics.components.snapshot.reserve |
| statistics.components.snapshot.used | string | query | False | Filter by statistics.components.snapshot.used |
| statistics.components.latency.raw | string | query | False | Filter by statistics.components.latency.raw |
| statistics.components.latency.average | string | query | False | Filter by statistics.components.latency.average |
| statistics.components.statistics_incomplete | string | query | False | Filter by statistics.components.statistics_incomplete |
| statistics.components.shared_storage_pool | string | query | False | Filter by statistics.components.shared_storage_pool |
| smart_container | string | query | False | Filter by smart_container<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[application] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        },
        "snapshots": {
          "href": "/api/resourcelink"
        }
      },
      "creation_timestamp": "string",
      "delete_data": true,
      "generation": 0,
      "name": "string",
      "nas": {
        "application_components": [
          {}
        ],
        "cifs_access": [
          {
            "access": "string",
            "user_or_group": "string"
          }
        ],
        "cifs_share_name": "string",
        "exclude_aggregates": [
          {
            "name": "string",
            "uuid": "string"
          }
        ],
        "nfs_access": [
          {
            "access": "string",
```

```
            "host": "string"
          }
        ],
        "protection_type": {
          "local_policy": "string",
          "local_rpo": "string",
          "remote_rpo": "string"
        }
      },
      "nvme": {
        "components": [
          {
            "name": "string",
            "os_type": "string",
            "performance": {
              "storage_service": {
                "name": "string"
              }
            },
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            },
            "subsystem": {
              "hosts": [
                {
                  "dh_hmac_chap": {
                    "group_size": "string",
                    "hash_function": "string"
                  },
                  "nqn": "string",
                  "priority": "string"
                }
              ],
              "name": "string",
              "os_type": "string",
              "uuid": "string"
            },
            "tiering": {
              "policy": "string"
            }
          }
        ],
        "os_type": "string",
```

```
        "rpo": {
          "local": {
            "name": "string",
            "policy": "string"
          },
          "remote": {
            "name": "string"
          }
        }
      },
      "protection_granularity": "string",
      "rpo": {
        "components": [
          {
            "name": "string",
            "rpo": {
              "local": {
                "description": "string",
                "name": "string"
              },
              "remote": {
                "description": "string",
                "name": "string"
              }
            },
            "uuid": "string"
          }
        ],
        "local": {
          "description": "string",
          "name": "string"
        },
        "remote": {
          "description": "string",
          "name": "string"
        }
      },
      "s3_bucket": {
        "application_components": [
          {
            "access_policies": [
              {
                "actions": [
                  "string"
                ],
                "conditions": [
```

```
                {
                  "delimiters": [
                    "string"
                  ],
                  "max_keys": [
                    "integer"
                  ],
                  "operator": "string",
                  "prefixes": [
                    "string"
                  ],
                  "source_ips": [
                    "string"
                  ],
                  "usernames": [
                    "string"
                  ]
                }
              ],
              "effect": "string",
              "principals": [
                "string"
              ],
              "resources": [
                "string"
              ],
              "sid": "string"
            }
          ],
          "bucket_endpoint_type": "string",
          "comment": "string",
          "default_retention_period": "string",
          "exclude_aggregates": [
            {
              "name": "string",
              "uuid": "string"
            }
          ],
          "name": "string",
          "nas_path": "string",
          "qos": {
            "policy": {
              "name": "string",
              "uuid": "string"
            }
          },
```

```
            "retention_mode": "string",
            "snapshot_policy": {
              "name": "string",
              "uuid": "string"
            },
            "storage_service": {
              "name": "string"
            },
            "uuid": "string",
            "versioning_state": "string"
          }
        ],
        "protection_type": {
          "remote_rpo": "string"
        }
      },
      "san": {
        "application_components": [
          {
            "igroup_name": "string",
            "name": "string",
            "os_type": "string",
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            },
            "storage_service": {
              "name": "string"
            },
            "tiering": {
              "policy": "string"
            }
          }
        ],
        "exclude_aggregates": [
          {
            "name": "string",
            "uuid": "string"
          }
        ],
        "os_type": "string",
        "protection_type": {
          "local_policy": "string",
          "local_rpo": "string",
```

```
          "remote_rpo": "string"
        }
      },
      "state": "string",
      "statistics": {
        "components": [
          {
            "iops": {
              "per_tb": 0,
              "total": 0
            },
            "latency": {
              "average": 0,
              "raw": 0
            },
            "name": "string",
            "snapshot": {
              "reserve": 0,
              "used": 0
            },
            "space": {
              "available": 0,
              "logical_used": 0,
              "provisioned": 0,
              "reserved_unused": 0,
              "savings": 0,
              "used": 0,
              "used_excluding_reserves": 0,
              "used_percent": 0
            },
            "storage_service": {
              "name": "string",
              "uuid": "string"
            },
            "uuid": "string"
          }
        ],
        "iops": {
          "per_tb": 0,
          "total": 0
        },
        "latency": {
          "average": 0,
          "raw": 0
        },
        "snapshot": {
```

```
            "reserve": 0,
            "used": 0
          },
          "space": {
            "available": 0,
            "logical_used": 0,
            "provisioned": 0,
            "reserved_unused": 0,
            "savings": 0,
            "used": 0,
            "used_excluding_reserves": 0,
            "used_percent": 0
          }
        },
        "svm": {
          "name": "string",
          "uuid": "string"
        },
        "template": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "string",
          "protocol": "string",
          "version": 0
        },
        "uuid": "string"
      }
    ]
  }
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |
| snapshots | href | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| element_count | integer | The number of storage elements (LUNs for SAN) of the database to maintain. Must be an even number between 2 and 16. Odd numbers will be rounded up to the next even number within range. |
| replication_factor | integer | The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

initiator_objects

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

protection_type

| Name | Type | Description |
|---|---|---|
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

secondary_igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
|---|---|---|
| dataset | dataset | |
| os_type | string | The name of the host OS running the application. |
| primary_igroup_name | string | The initiator group for the primary. |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

| Name | Type | Description |
|---|---|---|
| id | integer | The ID of an existing NFS export policy. |

| Name | Type | Description |
|---|---|---|
| name | string | The name of an existing NFS export policy. |

component

| Name | Type | Description |
|---|---|---|
| name | string | Name of the source component. |

svm

| Name | Type | Description |
|---|---|---|
| name | string | Name of the source SVM. |

origin

| Name | Type | Description |
|---|---|---|
| component | component | |
| svm | svm | |

flexcache

| Name | Type | Description |
|---|---|---|
| dr_cache | boolean | Dr-cache is a FlexCache volume create time option that has the same flexgroup-msid as that of the origin of a FlexCache volume. By default, dr-cache is disabled. The flexgroup-msid of the FlexCache volume does not need to be same as that of the origin of a FlexCache volume. |
| origin | origin | |

policy

| Name | Type | Description |
|---|---|---|
| name | string | The name of an existing QoS policy. |
| uuid | string | The UUID of an existing QoS policy. Usage: <UUID> |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as quot;Plt;num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period. |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| retention | retention | |
| snaplock_type | string | The SnapLock type of the smart container. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

nas_application_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | |
| flexcache | flexcache | |
| name | string | The name of the application component. |
| qos | qos | |
| scale_out | boolean | Denotes a Flexgroup. |
| share_count | integer | The number of shares in the application component. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Indicates whether Snapshot copy locking is enabled on the volume. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The CIFS access granted to the user or group. |
| user_or_group | string | The name of the CIFS user or group that will be granted access. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the aggregate to exclude. Usage: <aggregate name> |
| uuid | string | The ID of the aggregate to exclude. Usage: <UUID> |

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The NFS access granted. |
| host | string | The name of the NFS entity granted access. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

nas

A generic NAS application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| cifs_access | array[app_cifs_access] | The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access. |

| Name | Type | Description |
| --- | --- | --- |
| cifs_share_name | string | The name of the CIFS share. Usage: <Share> |
| exclude_aggregates | array[exclude_aggregates] | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

performance

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| qos | qos | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |

| Name | Type | Description |
|---|---|---|
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
|---|---|---|
| name | string | The local RPO of the application. |
| policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |

remote

| Name | Type | Description |
|---|---|---|
| name | string | The remote RPO of the application. |

rpo

| Name | Type | Description |
|---|---|---|
| local | local | |
| remote | remote | |

zapp_nvme

An NVME application.

| Name | Type | Description |
|---|---|---|
| components | array[components] | |
| os_type | string | The name of the host OS running the application. |
| rpo | rpo | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| mirrored | boolean | Specifies whether the redo log group should be mirrored. |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| ora_home | ora_home | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| copies | integer | The number of CRS volumes. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
|---|---|---|
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

local

| Name | Type | Description |
|---|---|---|
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the component. This indicates how often component Snapshot copies are automatically created. |

remote

| Name | Type | Description |
|---|---|---|
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | Component Name. |
| rpo | rpo | |
| uuid | string | Component UUID. |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the application. This indicates how often application Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the application. A remote RPO of zero indicates that the application is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| is_supported | boolean | Is RPO supported for this application? Generation 1 applications did not support Snapshot copies or MetroCluster. |
| local | local | |
| remote | remote | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| operator | string | Policy Condition Operator. |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[zapp_s3_bucket_application_components_access_policies_conditions] | conditions. |
| effect | string | Allow or Deny Access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Statement Identifier Usage: <(size 1..256)> |

snapshot_policy

| Name | Type | Description |
|---|---|---|
| name | string | The name of the snapshot policy that is used for the S3 bucket. |
| uuid | string | The UUID of the snapshot policy that is used for the S3 bucket. Usage: <UUID> |

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
| --- | --- | --- |
| access_policies | array[zapp_s3_bucket_application_components_access_policies] | The list of S3 objectstore policies to be created. |
| bucket_endpoint_type | string | The type of bucket. |
| capacity_tier | boolean | Prefer lower latency storage under similar media costs. |
| comment | string | Object Store Server Bucket Description Usage: <(size 1..256)> |
| default_retention_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as quot;Plt;num>Y" and quot;Plt;num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported. Usage: {{<integer> days\|years} \| none} |
| exclude_aggregates | array[exclude_aggregates] | |
| name | string | The name of the application component. |
| nas_path | string | The path to which the bucket corresponds to. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| retention_mode | string | The lock mode of the bucket. <br>compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects. <br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. For s3 type buckets, the default value is no_lock. |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| snapshot_policy | snapshot_policy | |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |
| versioning_state | string | Bucket Versioning State. For nas type buckets, this field is not set. For s3 type buckets, the default value is disabled. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| remote_rpo | string | The remote RPO of the application. |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |
| protection_type | protection_type | |

san_application_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
|---|---|---|
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| qos | qos | |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

san

A generic SAN application.

| Name | Type | Description |
|---|---|---|
| application_components | array[application_components] | |
| exclude_aggregates | array[exclude_aggregates] | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

access

| Name | Type | Description |
| --- | --- | --- |
| installer | string | SQL installer admin user name. |
| service_account | string | SQL service account user name. |

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| access | access | |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application component. |
| total | integer | The total number of IOPS being used by the application component. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this component. |
| raw | integer | The cumulative response time in microseconds for this component. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| reserve | integer | The amount of space reserved by the system for Snapshot copies. |
| used | integer | The amount of spacing currently in use by the system to store Snapshot copies. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application component. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application component. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service name. AFF systems support the extreme storage service. All other systems only support value. |
| uuid | string | The storage service UUID. |

components

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| latency | latency | |
| name | string | Component Name. |
| shared_storage_pool | boolean | An application component is considered to use a shared storage pool if storage elements for for other components reside on the same aggregate as storage elements for this component. |
| snapshot | snapshot | |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| statistics_incomplete | boolean | If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available. |
| storage_service | storage_service | |
| uuid | string | Component UUID. |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application. |
| total | integer | The total number of IOPS being used by the application. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this application. |
| raw | integer | The cumulative response time in microseconds for this application. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |
| used | integer | The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |

| Name | Type | Description |
| --- | --- | --- |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application. |

statistics

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| iops | iops | |
| latency | latency | |
| shared_storage_pool | boolean | An application is considered to use a shared storage pool if storage elements for multiple components reside on the same aggregate. |
| snapshot | snapshot | |
| space | space | |
| statistics_incomplete | boolean | If not all storage elements of the application are currently available, the returned statistics might only include data from those elements that were available. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | SVM Name. Either the SVM name or UUID must be provided to create an application. |
| uuid | string | SVM UUID. Either the SVM name or UUID must be provided to create an application. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

template

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The name of the template that was used to provision this application. |
| protocol | string | The protocol access of the template that was used to provision this application. |
| version | integer | The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| size | integer | The size of the desktops. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

hyper_v_access

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | Hyper-V service account. |

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of datastores to support. |
| size | integer | The size of the datastore. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| hypervisor | string | The name of the hypervisor hosting the application. |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

application

Applications

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Application Name. This field is user supplied when the application is created. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |

| Name | Type | Description |
| --- | --- | --- |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| smart_container | boolean | Identifies if this is a smart container or not. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| svm | svm | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |

| Name | Type | Description |
|---|---|---|
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Create an application

POST `/application/applications`

**Introduced In:** 9.6

Creates an application.

### Supported templates

- MongoDBOnSAN
- MultiComponentNAS
- NAS
- NVME
- OracleOnNFS
- OracleOnSAN

- OracleRACOnNFS
- OracleRACOnSAN
- S3
- SAN
- SAPHANAOnNFS
- SQLOnSAN
- SQLOnSMB
- VDIOnNAS
- VDIOnSAN
- VSIOnNAS
- VSIOnSAN

#### Template properties

The application APIs appear to be complex and long in this documentation because we document every possible template, of which there are currently 14. When creating an application, only a single template is used, so it is best to focus only on the template of interest. Other than the properties for the chosen template, only the `name` and `svm` of the application must be provided. The following three sections provided guidelines on using the properties of the templates, but the whole idea behind the templates is to automatically follow the best practices of the given application, so the only way to determine the exact list of required properties and default values is to dig in to the model section of the template. The templates are all top level properties of the application object with names matching the values returned by GET /application/templates .

#### Required properties

- `svm.uuid` or `svm.name` - The existing SVM in which to create the application.
- `name` - The name for the application.
- `<template>` - Properties for one template must be provided. In general, the following properties are required, however the naming of these may vary slightly from template to template.
  - `<template>.name` - The generic templates require names for the components of the application. Other templates name the components automatically.
  - `<template>.application_component.size` - This generally refers to the size of an application component, which may be spread across multiple underlying storage objects (volumes, LUNs, etc…).
  - One of the following must be specified:
- `nas.nfs_access` or an identifier (name or id) of an existing `export-policy`.
- `nas.cifs_access`
- `san.application_application.igroup_name`
  - To specify that a NAS application is not to be exposed via NFS nor CIFS:
- `nas.application_component.export_policy.name` is 'none', and
- `nas.application.cifs_share_name` is 'none'.
  - The name of the CIFS share can be provided through the `nas.application.cifs_share_name` attribute. If not provided, the CIFS share name will be the same as the `nas.application_component.name` attribute, including any suffix applied due to creating multiple

application components in one post.

  - `san.os_type` - All SAN applications require an os_type to be specified in some way. Some templates refer to this as the `hypervisor`.

#### Recommended optional properties

- `<template>` - The following properties are available in some templates.
  - `san.new_igroups.*` - SAN applications can use existing initiator groups or create new ones. When creating new initiator groups, `new_igroups.name` is required and the other properties may be used to fully specify the new initiator group.

#### Default property values

If not specified in POST, the follow default property values are assigned. It is recommended that most of these properties be provided explicitly rather than relying upon the defaults. The defaults are intended to make it as easy as possible to provision and connect to an application.

- `template.name` - Defaults to match the `<template>` provided. If specified, the value of this property must match the provided template properties.
- `<template>` - The majority of template properties have default values. The defaults may vary from template to template. See the model of each template for complete details. In general the following patterns are common across all template properties. The location of these properties varies from template to template.
  - `<template>.storage_service.name` - *value*
  - `<template>.protection_type.local_rpo` - *hourly* (Hourly snapshots)
  - `<template>.protection_type.remote_rpo` - *none* (Not MetroCluster)
  - `san.new_igroups.os_type` - Defaults to match the `os_type` provided for the application, but may need to be provided explicitly when using virtualization.

#### Optional components

A common pattern across many templates are objects that are optional, but once any property in the object is specified, other properties within the object become required. Many applications have optional components. For example, provisioning a database without a component to store the logs is supported. If the properties related to the logs are omitted, no storage will be provisioned for logs. But when the additional component is desired, the size is required. Specifying any other property of a component without specifying the size is not supported. In the model of each template, the required components are indicated with a red '*'. When a `size` property is listed as optional, that means the component itself is optional, and the size should be specified to include that component in the application.

#### POST body examples

1. Create a generic SAN application that exposes four LUNs to an existing initiator group, *igroup_1*.

```
  {
    "name": "app1",
    "svm": { "name": "svm1" },
    "san": {
      "os_type": "linux",
      "application_components": [
        { "name": "component1", "total_size": "10GB", "lun_count": 4,
"igroup_name": "igroup_1" }
      ]
    }
  }
```

1. Create an SQL application that can be accessed via initiator *iqn.2017-01.com.example:foo* from a new initiator group, *igroup_2*.

```
  {
    "name": "app2",
    "svm": { "name": "svm1" },
    "sql_on_san": {
      "db": { "size": "5GB" },
      "log": { "size": "1GB" },
      "temp_db": { "size": "2GB" },
      "igroup_name": "igroup_2",
      "new_igroups": [
        { "name": "igroup_2", "initiators": [ "iqn.2017-
01.com.example:foo" ] }
      ]
    }
  }
```

1. The following body creates the exact same SQL application, but manually provides all the defaults that were excluded from the previous call.

> The model of a *sql_on_san* application documents all these default values.

```
  {
    "name": "app3",
    "svm": { "name": "svm1" },
    "template": { "name": "sql_on_san" },
    "sql_on_san": {
      "os_type": "windows_2008",
      "server_cores_count": 8,
      "db": { "size": "5GB", "storage_service": { "name": "value" } },
      "log": { "size": "1GB", "storage_service": { "name": "value" } },
      "temp_db": { "size": "2GB", "storage_service": { "name": "value" }
},
      "igroup_name": "igroup_2",
      "new_igroups": [
        {
          "name": "igroup_2",
          "protocol": "mixed",
          "os_type": "windows",
          "initiators": [ "iqn.a.new.initiator" ]
        }
      ],
      "protection_type": { "local_rpo": "none" }
    }
  }
```

### Learn more

- DOC /application
- Asynchronous operations

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |

| Name | Type | Description |
| --- | --- | --- |
| delete_data | boolean | Should application storage elements be deleted? An application is considered to use storage elements from a shared storage pool. Possible values are 'true' and 'false'. If the value is 'true', the application will be deleted in its entirety. If the value is 'false', the storage elements will be disassociated from the application and preserved. The application will then be deleted. |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Application Name. This field is user supplied when the application is created. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |

| Name | Type | Description |
| --- | --- | --- |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| smart_container | boolean | Identifies if this is a smart container or not. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| svm | svm | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |

| Name | Type | Description |
| --- | --- | --- |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

**Example request**

```
{
  "creation_timestamp": "string",
  "generation": 0,
  "name": "string",
  "nas": {
    "application_components": [
      {}
    ],
    "cifs_access": [
      {
        "access": "string",
        "user_or_group": "string"
      }
    ],
    "cifs_share_name": "string",
    "exclude_aggregates": [
      {
        "name": "string",
        "uuid": "string"
      }
    ],
    "nfs_access": [
      {
        "access": "string",
        "host": "string"
      }
    ],
    "protection_type": {
      "local_policy": "string",
      "local_rpo": "string",
      "remote_rpo": "string"
    }
  },
  "nvme": {
    "components": [
      {
        "name": "string",
        "os_type": "string",
        "performance": {
          "storage_service": {
            "name": "string"
          }
        },
        "qos": {
```

```
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "subsystem": {
          "hosts": [
            {
              "dh_hmac_chap": {
                "controller_secret_key": "string",
                "group_size": "string",
                "hash_function": "string",
                "host_secret_key": "string"
              },
              "nqn": "string",
              "priority": "string"
            }
          ],
          "name": "string",
          "os_type": "string",
          "uuid": "string"
        },
        "tiering": {
          "control": "string",
          "object_stores": [
            {
              "name": "string"
            }
          ],
          "policy": "string"
        }
      }
    ],
    "os_type": "string",
    "rpo": {
      "local": {
        "name": "string",
        "policy": "string"
      },
      "remote": {
        "name": "string"
      }
    }
  },
  "protection_granularity": "string",
  "rpo": {
```

```
    "components": [
      {
        "name": "string",
        "rpo": {
          "local": {
            "description": "string",
            "name": "string"
          },
          "remote": {
            "description": "string",
            "name": "string"
          }
        },
        "uuid": "string"
      }
    ],
    "local": {
      "description": "string",
      "name": "string"
    },
    "remote": {
      "description": "string",
      "name": "string"
    }
  },
  "s3_bucket": {
    "application_components": [
      {
        "access_policies": [
          {
            "actions": [
              "string"
            ],
            "conditions": [
              {
                "delimiters": [
                  "string"
                ],
                "max_keys": [
                  "integer"
                ],
                "operator": "string",
                "prefixes": [
                  "string"
                ],
                "source_ips": [
```

```
                "string"
              ],
              "usernames": [
                "string"
              ]
            }
          ],
          "effect": "string",
          "principals": [
            "string"
          ],
          "resources": [
            "string"
          ],
          "sid": "string"
        }
      ],
      "bucket_endpoint_type": "string",
      "comment": "string",
      "default_retention_period": "string",
      "exclude_aggregates": [
        {
          "name": "string",
          "uuid": "string"
        }
      ],
      "name": "string",
      "nas_path": "string",
      "qos": {
        "policy": {
          "name": "string",
          "uuid": "string"
        }
      },
      "retention_mode": "string",
      "snapshot_policy": {
        "name": "string",
        "uuid": "string"
      },
      "storage_service": {
        "name": "string"
      },
      "uuid": "string",
      "versioning_state": "string"
    }
  ],
```

```
      "protection_type": {
        "remote_rpo": "string"
      }
    },
    "san": {
      "application_components": [
        {
          "igroup_name": "string",
          "name": "string",
          "os_type": "string",
          "qos": {
            "policy": {
              "name": "string",
              "uuid": "string"
            }
          },
          "storage_service": {
            "name": "string"
          },
          "tiering": {
            "control": "string",
            "object_stores": [
              {
                "name": "string"
              }
            ],
            "policy": "string"
          }
        }
      ],
      "exclude_aggregates": [
        {
          "name": "string",
          "uuid": "string"
        }
      ],
      "new_igroups": [
        {
          "comment": "string",
          "igroups": [
            {
              "name": "string",
              "uuid": "string"
            }
          ],
          "initiator_objects": [
```

```
          {
            "comment": "string",
            "name": "string"
          }
        ],
        "initiators": [
          "string"
        ],
        "name": "string",
        "os_type": "string",
        "protocol": "string"
      }
    ],
    "os_type": "string",
    "protection_type": {
      "local_policy": "string",
      "local_rpo": "string",
      "remote_rpo": "string"
    }
  },
  "state": "string",
  "statistics": {
    "components": [
      {
        "iops": {
          "per_tb": 0,
          "total": 0
        },
        "latency": {
          "average": 0,
          "raw": 0
        },
        "name": "string",
        "snapshot": {
          "reserve": 0,
          "used": 0
        },
        "space": {
          "available": 0,
          "logical_used": 0,
          "provisioned": 0,
          "reserved_unused": 0,
          "savings": 0,
          "used": 0,
          "used_excluding_reserves": 0,
          "used_percent": 0
```

```
        },
        "storage_service": {
          "name": "string",
          "uuid": "string"
        },
        "uuid": "string"
      }
    ],
    "iops": {
      "per_tb": 0,
      "total": 0
    },
    "latency": {
      "average": 0,
      "raw": 0
    },
    "snapshot": {
      "reserve": 0,
      "used": 0
    },
    "space": {
      "available": 0,
      "logical_used": 0,
      "provisioned": 0,
      "reserved_unused": 0,
      "savings": 0,
      "used": 0,
      "used_excluding_reserves": 0,
      "used_percent": 0
    }
  },
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "template": {
    "name": "string",
    "protocol": "string",
    "version": 0
  },
  "uuid": "string"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

### Response

```
Status: 201, Created
```

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65995775 | The size provided is too small. |
| 65995776 | The size provided is too large. |
| 92406300 | The "-capacity-tier" parameter cannot be set to "true" because tiering is not supported on the platform. |
| 92406301 | The "-storage_service" parameter cannot be specified because it is not supported on the platform." |

Also see the table of common errors in the Response body overview section of this documentation.

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| element_count | integer | The number of storage elements (LUNs for SAN) of the database to maintain. Must be an even number between 2 and 16. Odd numbers will be rounded up to the next even number within range. |
| replication_factor | integer | The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of an igroup to nest within a parent igroup. Mutually exclusive with initiators and initiator_objects. |
| uuid | string | The UUID of an igroup to nest within a parent igroup Usage: <UUID> |

initiator_objects

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| name | string | The WWPN, IQN, or Alias of the initiator. Mutually exclusive with nested igroups and the initiators array. |

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

protection_type

| Name | Type | Description |
|---|---|---|
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

secondary_igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
| --- | --- | --- |
| dataset | dataset | |
| new_igroups | array[mongo_db_on_san_new_igroups] | The list of initiator groups to create. |
| os_type | string | The name of the host OS running the application. |
| primary_igroup_name | string | The initiator group for the primary. |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ID of an existing NFS export policy. |
| name | string | The name of an existing NFS export policy. |

component

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source component. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source SVM. |

origin

| Name | Type | Description |
| --- | --- | --- |
| component | component | |
| svm | svm | |

flexcache

| Name | Type | Description |
| --- | --- | --- |
| dr_cache | boolean | Dr-cache is a FlexCache volume create time option that has the same flexgroup-msid as that of the origin of a FlexCache volume. By default, dr-cache is disabled. The flexgroup-msid of the FlexCache volume does not need to be same as that of the origin of a FlexCache volume. |
| origin | origin | |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an existing QoS policy. |
| uuid | string | The UUID of an existing QoS policy. Usage: <UUID> |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as quot;Plt;num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period. |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| retention | retention | |
| snaplock_type | string | The SnapLock type of the smart container. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object-store to use. |

nas_application_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the container(s) |
| object_stores | array[object_stores] | |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | |
| flexcache | flexcache | |
| name | string | The name of the application component. |
| qos | qos | |
| scale_out | boolean | Denotes a Flexgroup. |
| share_count | integer | The number of shares in the application component. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Indicates whether Snapshot copy locking is enabled on the volume. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

| Name | Type | Description |
|---|---|---|
| access | string | The CIFS access granted to the user or group. |
| user_or_group | string | The name of the CIFS user or group that will be granted access. |

exclude_aggregates

| Name | Type | Description |
|---|---|---|
| name | string | The name of the aggregate to exclude. Usage: <aggregate name> |
| uuid | string | The ID of the aggregate to exclude. Usage: <UUID> |

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

| Name | Type | Description |
|---|---|---|
| access | string | The NFS access granted. |
| host | string | The name of the NFS entity granted access. |

protection_type

| Name | Type | Description |
|---|---|---|
| local_policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

nas

A generic NAS application.

| Name | Type | Description |
|---|---|---|
| application_components | array[application_components] | |
| cifs_access | array[app_cifs_access] | The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access. |
| cifs_share_name | string | The name of the CIFS share. Usage: <Share> |
| exclude_aggregates | array[exclude_aggregates] | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

performance

| Name | Type | Description |
|---|---|---|
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
|---|---|---|
| controller_secret_key | string | Authentication Controller Secret. |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |
| host_secret_key | string | Authentication Host Secret. |

hosts

| Name | Type | Description |
|---|---|---|
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
|---|---|---|
| hosts | array[hosts] | |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| control | string | Storage tiering placement rules for the container(s) |
| object_stores | array[object_stores] | |
| policy | string | The storage tiering type of the application component. |

components

| Name | Type | Description |
|---|---|---|
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| qos | qos | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local RPO of the application. |
| policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |

remote

| Name | Type | Description |
| --- | --- | --- |
| name | string | The remote RPO of the application. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| local | local | |
| remote | remote | |

zapp_nvme

An NVME application.

| Name | Type | Description |
|---|---|---|
| components | array[components] | |
| os_type | string | The name of the host OS running the application. |
| rpo | rpo | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| mirrored | boolean | Specifies whether the redo log group should be mirrored. |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[oracle_on_san_new_igroups] | The list of initiator groups to create. |
| ora_home | ora_home | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| copies | integer | The number of CRS volumes. |
| size | integer | The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |

| Name | Type | Description |
| --- | --- | --- |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| new_igroups | array[oracle_rac_on_san_new_igroups] | The list of initiator groups to create. |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the component. This indicates how often component Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | Component Name. |
| rpo | rpo | |
| uuid | string | Component UUID. |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the application. This indicates how often application Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the application. A remote RPO of zero indicates that the application is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| is_supported | boolean | Is RPO supported for this application? Generation 1 applications did not support Snapshot copies or MetroCluster. |
| local | local | |
| remote | remote | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| operator | string | Policy Condition Operator. |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[zapp_s3_bucket_application_components_access_policies_conditions] | conditions. |
| effect | string | Allow or Deny Access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Statement Identifier Usage: <(size 1..256)> |

snapshot_policy

| Name | Type | Description |
|---|---|---|
| name | string | The name of the snapshot policy that is used for the S3 bucket. |
| uuid | string | The UUID of the snapshot policy that is used for the S3 bucket. Usage: <UUID> |

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
| --- | --- | --- |
| access_policies | array[zapp_s3_bucket_application_components_access_policies] | The list of S3 objectstore policies to be created. |
| bucket_endpoint_type | string | The type of bucket. |
| capacity_tier | boolean | Prefer lower latency storage under similar media costs. |
| comment | string | Object Store Server Bucket Description Usage: <(size 1..256)> |
| default_retention_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as quot;Plt;num>Y" and quot;Plt;num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported. Usage: {{<integer> days\|years} \| none} |
| exclude_aggregates | array[exclude_aggregates] | |
| name | string | The name of the application component. |
| nas_path | string | The path to which the bucket corresponds to. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| retention_mode | string | The lock mode of the bucket. <br>compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects. <br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. For s3 type buckets, the default value is no_lock. |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| snapshot_policy | snapshot_policy | |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |
| versioning_state | string | Bucket Versioning State. For nas type buckets, this field is not set. For s3 type buckets, the default value is disabled. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| remote_rpo | string | The remote RPO of the application. |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |
| protection_type | protection_type | |

san_application_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| control | string | Storage tiering placement rules for the container(s) |
| object_stores | array[object_stores] | |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
|---|---|---|
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| qos | qos | |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

san

A generic SAN application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| exclude_aggregates | array[exclude_aggregates] | |
| new_igroups | array[san_new_igroups] | The list of initiator groups to create. |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| new_igroups | array[sql_on_san_new_igroups] | The list of initiator groups to create. |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

access

| Name | Type | Description |
| --- | --- | --- |
| installer | string | SQL installer admin user name. |

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | SQL service account user name. |

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| access | access | |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application component. |
| total | integer | The total number of IOPS being used by the application component. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this component. |
| raw | integer | The cumulative response time in microseconds for this component. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| reserve | integer | The amount of space reserved by the system for Snapshot copies. |
| used | integer | The amount of spacing currently in use by the system to store Snapshot copies. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application component. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application component. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service name. AFF systems support the extreme storage service. All other systems only support value. |
| uuid | string | The storage service UUID. |

components

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| latency | latency | |
| name | string | Component Name. |
| shared_storage_pool | boolean | An application component is considered to use a shared storage pool if storage elements for for other components reside on the same aggregate as storage elements for this component. |
| snapshot | snapshot | |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| statistics_incomplete | boolean | If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available. |
| storage_service | storage_service | |
| uuid | string | Component UUID. |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application. |
| total | integer | The total number of IOPS being used by the application. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this application. |
| raw | integer | The cumulative response time in microseconds for this application. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |
| used | integer | The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |

| Name | Type | Description |
| --- | --- | --- |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application. |

statistics

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| iops | iops | |
| latency | latency | |
| shared_storage_pool | boolean | An application is considered to use a shared storage pool if storage elements for multiple components reside on the same aggregate. |
| snapshot | snapshot | |
| space | space | |
| statistics_incomplete | boolean | If not all storage elements of the application are currently available, the returned statistics might only include data from those elements that were available. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | SVM Name. Either the SVM name or UUID must be provided to create an application. |
| uuid | string | SVM UUID. Either the SVM name or UUID must be provided to create an application. |

self_link

template

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the template that was used to provision this application. |
| protocol | string | The protocol access of the template that was used to provision this application. |
| version | integer | The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| size | integer | The size of the desktops. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

hyper_v_access

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | Hyper-V service account. |

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| protocol | string | The protocol of the new initiator group. |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hypervisor | string | The name of the hypervisor hosting the application. |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[vdi_on_san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of datastores to support. |
| size | integer | The size of the datastore. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| protocol | string | The protocol of the new initiator group. |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
|---|---|---|
| datastore | datastore | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[vsi_on_san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |

application

Applications

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |
| delete_data | boolean | Should application storage elements be deleted? An application is considered to use storage elements from a shared storage pool. Possible values are 'true' and 'false'. If the value is 'true', the application will be deleted in its entirety. If the value is 'false', the storage elements will be disassociated from the application and preserved. The application will then be deleted. |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Application Name. This field is user supplied when the application is created. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |

| Name | Type | Description |
| --- | --- | --- |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| smart_container | boolean | Identifies if this is a smart container or not. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| svm | svm | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |

| Name | Type | Description |
|---|---|---|
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve application components

GET `/application/applications/{application.uuid}/components`

**Introduced In:** 9.6

Retrieves application components.

### Overview

The application component object exposes how to access an application. Most application interfaces abstract away the underlying ONTAP storage elements, but this interface exposes what is necessary to connect to and uses the storage that is provisioned for an application. See the application component model for a detailed description of each property.

### Query examples

Queries are limited on this API. Most of the details are nested under the `nfs_access`, `cifs_access`, or `san_access` properties, but those properties do not support queries, and properties nested under those properties cannot be requested individually in the current release.

The following query returns all application components with names beginning in *secondary*.

```
GET
/application/applications/{application.uuid}/components?name=secondary*
```

The following query returns all application components at the *extreme* storage service.

```
GET
/application/applications/{application.uuid}/components?storage_service.na
me=extreme
```

### Learn more

- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | query | False | Filter by UUID |
| name | string | query | False | Filter by name |
| storage_service.name | string | query | False | Filter by storage_service.name |
| storage_service.uuid | string | query | False | Filter by storage_service.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[application_component] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "string"
      },
      "backing_storage": {
        "luns": [
          {
            "creation_timestamp": "string",
            "path": "string",
            "size": 0,
            "uuid": "string"
          }
        ],
        "namespaces": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "creation_timestamp": "string",
            "name": "string",
```

```
            "size": 0,
            "uuid": "string"
          }
        ],
        "volumes": [
          {
            "creation_timestamp": "string",
            "name": "string",
            "size": 0,
            "uuid": "string"
          }
        ]
      },
      "cifs_access": [
        {
          "backing_storage": {
            "type": "string",
            "uuid": "string"
          },
          "ips": [
            "string"
          ],
          "path": "string",
          "permissions": [
            {
              "access": "string",
              "user_or_group": "string"
            }
          ],
          "server": {
            "name": "string"
          },
          "share": {
            "name": "string"
          }
        }
      ],
      "file_system": "string",
      "host_management_url": "string",
      "host_name": "string",
      "name": "string",
      "nfs_access": [
        {
          "backing_storage": {
            "type": "string",
            "uuid": "string"
```

```
          },
          "export_policy": {
            "name": "string"
          },
          "ips": [
            "string"
          ],
          "path": "string",
          "permissions": [
            {
              "access": "string",
              "host": "string"
            }
          ]
        }
      ],
      "nvme_access": [
        {
          "backing_storage": {
            "type": "string",
            "uuid": "string"
          },
          "subsystem_map": {
            "anagrpid": "string",
            "nsid": "string",
            "subsystem": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "hosts": [
                {
                  "_links": {
                    "self": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    }
                  },
                  "nqn": "string"
                }
              ],
              "name": "string",
              "uuid": "string"
            }
```

```
          }
        }
      ],
      "protection_groups": [
        {
          "name": "string",
          "rpo": {
            "local": {
              "description": "string",
              "name": "string"
            },
            "remote": {
              "description": "string",
              "name": "string"
            }
          },
          "uuid": "string"
        }
      ],
      "san_access": [
        {
          "backing_storage": {
            "type": "string",
            "uuid": "string"
          },
          "lun_mappings": [
            {
              "fcp": [
                {
                  "interface": {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "fc_lif1",
                    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                    "wwpn": "20:00:00:50:56:b4:13:a8"
                  }
                }
              ],
              "igroup": {
                "initiators": [
                  "string"
                ],
                "name": "string",
```

```
                "uuid": "string"
              },
              "iscsi": [
                {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "interface": {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "ip": {
                      "address": "10.10.10.7"
                    },
                    "name": "lif1",
                    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                  },
                  "port": 3260
                }
              ],
              "lun_id": 0
            }
          ],
          "serial_number": "string"
        }
      ],
      "storage_service": {
        "name": "string",
        "uuid": "string"
      },
      "svm": {
        "name": "string",
        "uuid": "string"
      },
      "uuid": "string"
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application name |
| uuid | string | The application UUID. Valid in URL. |

application_lun_object

LUN object

| Name | Type | Description |
|---|---|---|
| creation_timestamp | string | LUN creation time |
| path | string | LUN path |
| size | integer | LUN size |
| uuid | string | LUN UUID |

application_namespace_object

Namespace object

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| creation_timestamp | string | Namespace creation time |
| name | string | Namespace name |
| size | integer | Namespace size |
| uuid | string | Namespace UUID |

application_volume_object

Volume object

| Name | Type | Description |
| --- | --- | --- |
| creation_timestamp | string | Creation time |
| name | string | Name |
| size | integer | Size |
| uuid | string | UUID |

application_backing_storage

| Name | Type | Description |
| --- | --- | --- |
| luns | array[application_lun_object] | |
| namespaces | array[application_namespace_object] | |
| volumes | array[application_volume_object] | |

backing_storage

| Name | Type | Description |
| --- | --- | --- |
| type | string | Backing storage type |
| uuid | string | Backing storage UUID |

permissions

| Name | Type | Description |
| --- | --- | --- |
| access | string | Access granted to the user or group |
| user_or_group | string | User or group |

server

| Name | Type | Description |
| --- | --- | --- |
| name | string | Server name |

share

| Name | Type | Description |
| --- | --- | --- |
| name | string | Share name |

application_cifs_properties

| Name | Type | Description |
| --- | --- | --- |
| backing_storage | backing_storage | |
| ips | array[string] | |
| path | string | Junction path |
| permissions | array[permissions] | |
| server | server | |
| share | share | |

export_policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | Export policy name |

permissions

| Name | Type | Description |
| --- | --- | --- |
| access | string | Access granted to the host |
| host | string | Host granted access |

application_nfs_properties

| Name | Type | Description |
|---|---|---|
| backing_storage | backing_storage | |
| export_policy | export_policy | |
| ips | array[string] | |
| path | string | Junction path |
| permissions | array[permissions] | |

self

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

hosts

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| nqn | string | Host |

subsystem

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| hosts | array[hosts] | |
| name | string | Subsystem name |
| uuid | string | Subsystem UUID |

subsystem_map

Subsystem map object

| Name | Type | Description |
|---|---|---|
| anagrpid | string | Subsystem ANA group ID |
| nsid | string | Subsystem namespace ID |

| Name | Type | Description |
| --- | --- | --- |
| subsystem | subsystem | |

application_nvme_access

Application NVME access

| Name | Type | Description |
| --- | --- | --- |
| backing_storage | backing_storage | |
| is_clone | boolean | Clone |
| subsystem_map | subsystem_map | Subsystem map object |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This includes details on the snapshot schedule. |
| name | string | The local RPO of the component. This indicates how often component snapshots are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| local | local | |
| remote | remote | |

application_protection_groups

| Name | Type | Description |
|---|---|---|
| name | string | Protection group name |
| rpo | rpo | |
| uuid | string | Protection group UUID |

fc_interface_reference

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

application_san_access_fcp_endpoint

A Fibre Channel Protocol (FCP) access endpoint for the LUN.

| Name | Type | Description |
|---|---|---|
| interface | fc_interface_reference | An FC interface. |

igroup

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |
| name | string | Igroup name |
| uuid | string | Igroup UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

application_san_access_iscsi_endpoint

An iSCSI access endpoint for the LUN.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| port | integer | The TCP port number of the iSCSI access endpoint. |

application_lun_mapping_object

| Name | Type | Description |
|---|---|---|
| fcp | array[application_san_access_fcp_endpoint] | All possible Fibre Channel Protocol (FCP) access endpoints for the LUN. |
| igroup | igroup | |
| iscsi | array[application_san_access_iscsi_endpoint] | All possible iSCSI access endpoints for the LUN. |
| lun_id | integer | LUN ID |

application_san_access

| Name | Type | Description |
|---|---|---|
| backing_storage | backing_storage | |
| is_clone | boolean | Clone |
| lun_mappings | array[application_lun_mapping_object] | |
| serial_number | string | LUN serial number |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | Storage service name |
| uuid | string | Storage service UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM name |
| uuid | string | SVM UUID |

application_component

Application component

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | application | |
| backing_storage | application_backing_storage | |
| cifs_access | array[application_cifs_properties] | |
| file_system | string | Defines the type of file system that will be installed on this application component. |
| host_management_url | string | Host management URL |
| host_name | string | L2 Host FQDN |
| name | string | Application component name |

| Name | Type | Description |
|---|---|---|
| nfs_access | array[application_nfs_properties] | |
| nvme_access | array[application_nvme_access] | |
| protection_groups | array[application_protection_groups] | |
| san_access | array[application_san_access] | |
| storage_service | storage_service | |
| svm | svm | |
| uuid | string | The application component UUID. Valid in URL. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve application component snapshots

GET `/application/applications/{application.uuid}/components/{component.uuid}/snapshots`

**Introduced In:** 9.6

Retrieves snapshots of an application component.

This endpoint is only supported for Maxdata template applications.

Component snapshots are essentially more granular application snapshots. There is no difference beyond the

scope of the operation.

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- GET /application/applications/{uuid}/snapshots
- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| component.uuid | string | path | True | Application Component UUID |
| component.name | string | query | False | Filter by Application Component Name |
| uuid | string | query | False | Filter by uuid |
| name | string | query | False | Filter by name |
| consistency_type | string | query | False | Filter by consistency_type |
| comment | string | query | False | Filter by comment |
| create_time | string | query | False | Filter by create_time |
| is_partial | string | query | False | Filter by is_partial |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[application_component_snapshot] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "string"
      },
      "comment": "string",
      "component": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "string"
      },
      "consistency_type": "string",
      "create_time": "string",
      "name": "string",
      "svm": {
        "name": "string",
        "uuid": "string"
      },
      "uuid": "string"
```

```
      }
    ]
  }
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application Name |
| uuid | string | Application UUID. Valid in URL |

component

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Component Name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM Name |
| uuid | string | SVM UUID |

application_component_snapshot

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | application | |
| comment | string | Comment. Valid in POST |
| component | component | |
| consistency_type | string | Consistency Type. This is for categorization only. A snapshot should not be set to application consistent unless the host application is quiesced for the snapshot. Valid in POST |
| create_time | string | Creation Time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | Snapshot name. Valid in POST |
| svm | svm | |
| uuid | string | Snapshot UUID. Valid in URL |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Create an application component snapshot

POST
`/application/applications/{application.uuid}/components/{component.uuid}/snapshots`

**Introduced In:** 9.6

Creates a snapshot of an application component.

This endpoint is only supported for Maxdata template applications.

### Required properties

- `name`

### Recommended optional properties

- `consistency_type` - Track whether this snapshot is *application* or *crash* consistent.
  Component snapshots are essentially more granular application snapshots. There is no difference beyond the scope of the operation.

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- GET /application/applications/{uuid}/snapshots
- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| component.uuid | string | path | True | Application Component UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| application | application | |
| comment | string | Comment. Valid in POST |
| component | component | |

| Name | Type | Description |
|---|---|---|
| consistency_type | string | Consistency Type. This is for categorization only. A snapshot should not be set to application consistent unless the host application is quiesced for the snapshot. Valid in POST |
| create_time | string | Creation Time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | Snapshot name. Valid in POST |
| svm | svm | |
| uuid | string | Snapshot UUID. Valid in URL |

**Example request**

```
{
  "application": {
    "name": "string",
    "uuid": "string"
  },
  "comment": "string",
  "component": {
    "name": "string",
    "uuid": "string"
  },
  "consistency_type": "string",
  "create_time": "string",
  "name": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

### Response

```
Status: 201, Created
```

### Error

```
Status: Default, Error
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

application

| Name | Type | Description |
|---|---|---|
| name | string | Application Name |
| uuid | string | Application UUID. Valid in URL |

component

| Name | Type | Description |
|---|---|---|
| name | string | Component Name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM Name |
| uuid | string | SVM UUID |

application_component_snapshot

| Name | Type | Description |
|---|---|---|
| application | application | |
| comment | string | Comment. Valid in POST |
| component | component | |
| consistency_type | string | Consistency Type. This is for categorization only. A snapshot should not be set to application consistent unless the host application is quiesced for the snapshot. Valid in POST |

| Name | Type | Description |
|---|---|---|
| create_time | string | Creation Time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | Snapshot name. Valid in POST |
| svm | svm | |
| uuid | string | Snapshot UUID. Valid in URL |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Delete an application component snapshot

DELETE

`/application/applications/{application.uuid}/components/{component.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Delete a snapshot of an application component.

This endpoint is only supported for Maxdata template applications.

Component snapshots are essentially more granular application snapshots. There is no difference beyond the scope of the operation.

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- DELETE /application/applications/{uuid}/snapshots
- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| component.uuid | string | path | True | Application Component UUID |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

### Response

```
Status: 202, Accepted
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

#### See Definitions

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve a snapshot for a specific application component

GET `/application/applications/{application.uuid}/components/{component.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Retrieve a snapshot of an application component.

This endpoint is only supported for Maxdata template applications.

Component snapshots are essentially more granular application snapshots. There is no difference beyond the scope of the operation.

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- GET /application/applications/{uuid}/snapshots
- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| component.uuid | string | path | True | Application Component UUID |
| uuid | string | path | True | Snapshot UUID |
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | application | |
| comment | string | Comment. Valid in POST |
| component | component | |

| Name | Type | Description |
| --- | --- | --- |
| consistency_type | string | Consistency Type. This is for categorization only. A snapshot should not be set to application consistent unless the host application is quiesced for the snapshot. Valid in POST |
| create_time | string | Creation Time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | Snapshot name. Valid in POST |
| svm | svm | |
| uuid | string | Snapshot UUID. Valid in URL |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "string"
  },
  "comment": "string",
  "component": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "string"
  },
  "consistency_type": "string",
  "create_time": "string",
  "name": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application Name |
| uuid | string | Application UUID. Valid in URL |

component

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Component Name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM Name |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Restore an application component snapshot

POST
`/application/applications/{application.uuid}/components/{component.uuid}/snapshots/{uuid}/restore`

**Introduced In:** 9.6

Restore a snapshot of an application component.

This endpoint is only supported for Maxdata template applications.

Component snapshots are essentially more granular application snapshots. There is no difference beyond the scope of the operation.

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- POST /application/applications/{application.uuid}/snapshots/{uuid}/restore
- DOC /application
- Asynchronous operations

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| component.uuid | string | path | True | Application Component UUID |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Response

```
Status: 201, Created
```

### Error

```
Status: Default, Error
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve an application component

GET `/application/applications/{application.uuid}/components/{uuid}`

**Introduced In:** 9.6

Retrieves an application component.

### Overview

The application component object exposes how to access an application. Most application interfaces abstract away the underlying ONTAP storage elements, but this interface exposes what is necessary to connect to and uses the storage that is provisioned for an application. See the application component model for a detailed description of each property.

### Access

Each application component can be accessed via NFS, CIFS, or SAN. NFS and CIFS access can be enabled simultaneously. Each access section includes a `backing_storage` property. This property is used to correlate the storage elements with the access elements of the application. The `backing_storage` portion of the access section provides the `type` and `uuid` of the backing storage. There is another `backing_storage` property at the same level as the access properties which contains lists of backing storage elements corresponding to the types listed in the access section.

### Learn more

- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | path | True | Application component UUID |
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | application | |
| backing_storage | application_backing_storage | |
| cifs_access | array[application_cifs_properties] | |
| file_system | string | Defines the type of file system that will be installed on this application component. |
| host_management_url | string | Host management URL |

| Name | Type | Description |
| --- | --- | --- |
| host_name | string | L2 Host FQDN |
| name | string | Application component name |
| nfs_access | array[application_nfs_properties] | |
| nvme_access | array[application_nvme_access] | |
| protection_groups | array[application_protection_groups] | |
| san_access | array[application_san_access] | |
| storage_service | storage_service | |
| svm | svm | |
| uuid | string | The application component UUID. Valid in URL. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "string"
  },
  "backing_storage": {
    "luns": [
      {
        "creation_timestamp": "string",
        "path": "string",
        "size": 0,
        "uuid": "string"
      }
    ],
    "namespaces": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "creation_timestamp": "string",
        "name": "string",
        "size": 0,
        "uuid": "string"
      }
    ],
    "volumes": [
      {
        "creation_timestamp": "string",
        "name": "string",
        "size": 0,
        "uuid": "string"
      }
```

```
      ]
    },
    "cifs_access": [
      {
        "backing_storage": {
          "type": "string",
          "uuid": "string"
        },
        "ips": [
          "string"
        ],
        "path": "string",
        "permissions": [
          {
            "access": "string",
            "user_or_group": "string"
          }
        ],
        "server": {
          "name": "string"
        },
        "share": {
          "name": "string"
        }
      }
    ],
    "file_system": "string",
    "host_management_url": "string",
    "host_name": "string",
    "name": "string",
    "nfs_access": [
      {
        "backing_storage": {
          "type": "string",
          "uuid": "string"
        },
        "export_policy": {
          "name": "string"
        },
        "ips": [
          "string"
        ],
        "path": "string",
        "permissions": [
          {
            "access": "string",
```

```
          "host": "string"
        }
      ]
    }
  ],
  "nvme_access": [
    {
      "backing_storage": {
        "type": "string",
        "uuid": "string"
      },
      "subsystem_map": {
        "anagrpid": "string",
        "nsid": "string",
        "subsystem": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "hosts": [
            {
              "_links": {
                "self": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                }
              },
              "nqn": "string"
            }
          ],
          "name": "string",
          "uuid": "string"
        }
      }
    }
  ],
  "protection_groups": [
    {
      "name": "string",
      "rpo": {
        "local": {
          "description": "string",
          "name": "string"
        },
```

```
        "remote": {
          "description": "string",
          "name": "string"
        }
      },
      "uuid": "string"
    }
  ],
  "san_access": [
    {
      "backing_storage": {
        "type": "string",
        "uuid": "string"
      },
      "lun_mappings": [
        {
          "fcp": [
            {
              "interface": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "fc_lif1",
                "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                "wwpn": "20:00:00:50:56:b4:13:a8"
              }
            }
          ],
          "igroup": {
            "initiators": [
              "string"
            ],
            "name": "string",
            "uuid": "string"
          },
          "iscsi": [
            {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "interface": {
                "_links": {
```

```
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "ip": {
                    "address": "10.10.10.7"
                  },
                  "name": "lif1",
                  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                },
                "port": 3260
              }
            ],
            "lun_id": 0
          }
        ],
        "serial_number": "string"
      }
    ],
    "storage_service": {
      "name": "string",
      "uuid": "string"
    },
    "svm": {
      "name": "string",
      "uuid": "string"
    },
    "uuid": "string"
  }
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application name |
| uuid | string | The application UUID. Valid in URL. |

application_lun_object

LUN object

| Name | Type | Description |
|---|---|---|
| creation_timestamp | string | LUN creation time |
| path | string | LUN path |
| size | integer | LUN size |
| uuid | string | LUN UUID |

application_namespace_object

Namespace object

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| creation_timestamp | string | Namespace creation time |
| name | string | Namespace name |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Namespace size |
| uuid | string | Namespace UUID |

application_volume_object

Volume object

| Name | Type | Description |
| --- | --- | --- |
| creation_timestamp | string | Creation time |
| name | string | Name |
| size | integer | Size |
| uuid | string | UUID |

application_backing_storage

| Name | Type | Description |
| --- | --- | --- |
| luns | array[application_lun_object] | |
| namespaces | array[application_namespace_object] | |
| volumes | array[application_volume_object] | |

backing_storage

| Name | Type | Description |
| --- | --- | --- |
| type | string | Backing storage type |
| uuid | string | Backing storage UUID |

permissions

| Name | Type | Description |
| --- | --- | --- |
| access | string | Access granted to the user or group |
| user_or_group | string | User or group |

server

| Name | Type | Description |
|---|---|---|
| name | string | Server name |

share

| Name | Type | Description |
|---|---|---|
| name | string | Share name |

application_cifs_properties

| Name | Type | Description |
|---|---|---|
| backing_storage | backing_storage | |
| ips | array[string] | |
| path | string | Junction path |
| permissions | array[permissions] | |
| server | server | |
| share | share | |

export_policy

| Name | Type | Description |
|---|---|---|
| name | string | Export policy name |

permissions

| Name | Type | Description |
|---|---|---|
| access | string | Access granted to the host |
| host | string | Host granted access |

application_nfs_properties

| Name | Type | Description |
|---|---|---|
| backing_storage | backing_storage | |
| export_policy | export_policy | |
| ips | array[string] | |
| path | string | Junction path |

| Name | Type | Description |
|---|---|---|
| permissions | array[permissions] | |

self

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

hosts

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| nqn | string | Host |

subsystem

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| hosts | array[hosts] | |
| name | string | Subsystem name |
| uuid | string | Subsystem UUID |

subsystem_map

Subsystem map object

| Name | Type | Description |
|---|---|---|
| anagrpid | string | Subsystem ANA group ID |
| nsid | string | Subsystem namespace ID |
| subsystem | subsystem | |

application_nvme_access

Application NVME access

| Name | Type | Description |
| --- | --- | --- |
| backing_storage | backing_storage | |
| is_clone | boolean | Clone |
| subsystem_map | subsystem_map | Subsystem map object |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This includes details on the snapshot schedule. |
| name | string | The local RPO of the component. This indicates how often component snapshots are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| local | local | |
| remote | remote | |

application_protection_groups

| Name | Type | Description |
| --- | --- | --- |
| name | string | Protection group name |
| rpo | rpo | |

| Name | Type | Description |
|---|---|---|
| uuid | string | Protection group UUID |

fc_interface_reference

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

application_san_access_fcp_endpoint

A Fibre Channel Protocol (FCP) access endpoint for the LUN.

| Name | Type | Description |
|---|---|---|
| interface | fc_interface_reference | An FC interface. |

igroup

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |
| name | string | Igroup name |
| uuid | string | Igroup UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

application_san_access_iscsi_endpoint

An iSCSI access endpoint for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| port | integer | The TCP port number of the iSCSI access endpoint. |

application_lun_mapping_object

| Name | Type | Description |
| --- | --- | --- |
| fcp | array[application_san_access_fcp_endpoint] | All possible Fibre Channel Protocol (FCP) access endpoints for the LUN. |
| igroup | igroup | |
| iscsi | array[application_san_access_iscsi_endpoint] | All possible iSCSI access endpoints for the LUN. |
| lun_id | integer | LUN ID |

application_san_access

| Name | Type | Description |
| --- | --- | --- |
| backing_storage | backing_storage | |

| Name | Type | Description |
| --- | --- | --- |
| is_clone | boolean | Clone |
| lun_mappings | array[application_lun_mapping_object] | |
| serial_number | string | LUN serial number |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name |
| uuid | string | Storage service UUID |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | SVM name |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage application snapshots

## Manage application snapshots

### Overview

Applications support snapshots across all member storage elements. These snapshots can be created and restored at any time or as scheduled. Most applications have hourly snapshots enabled by default, unless the RPO setting is overridden during the creation of the application. An application snapshot can be flagged as either *application consistent*, or *crash consistent*. From an ONTAP perspective, there is no difference between these two consistency types. These types are available for record keeping so that snapshots taken after the application is quiesced (application consistent) can be tracked separately from those snapshots taken without first quiescing the application (crash consistent). By default, all application snapshots are flagged to be *crash consistent*, and snapshots taken at a scheduled time are also considered *crash consistent*.

The functionality provided by these APIs is not integrated with the host application. Snapshots have limited value without host coordination, so the use of the SnapCenter Backup Management suite is recommended to ensure correct interaction between host applications and ONTAP.

## Retrieve an application snapshot

GET `/application/applications/{application.uuid}/snapshots`

**Introduced In:** 9.6

Retrieves snapshots of an application.

### Query examples

The following query returns all snapshots from May 4, 2017 EST. For readability, the colon (`:`) is left in this example. For an actual call, they should be escaped as `%3A`.

```
GET
/application/applications/{application.uuid}/snapshots?create_time=2017-
05-04T00:00:00-05:00..2017-05-04T23:59:59-05:00
```

The following query returns all snapshots that have been flagged as *application consistent*.

```
GET
/application/applications/{application.uuid}/snapshots?consistency_type=ap
plication
```

### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- DOC /application

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | query | False | Filter by UUID |
| name | string | query | False | Filter by name |
| consistency_type | string | query | False | Filter by consistency_type |
| components.name | string | query | False | Filter by components.name |
| components.uuid | string | query | False | Filter by components.uuid |
| comment | string | query | False | Filter by comment |
| create_time | string | query | False | Filter by create_time |
| is_partial | string | query | False | Filter by is_partial |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[application_snapshot] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "string"
      },
      "comment": "string",
      "components": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "string",
          "uuid": "string"
        }
      ],
      "consistency_type": "string",
      "create_time": "string",
      "name": "string",
      "svm": {
        "name": "string",
        "uuid": "string"
```

```
        },
        "uuid": "string"
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application name |
| uuid | string | The application UUID. Valid in URL. |

components

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Component name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM name |
| uuid | string | SVM UUID |

application_snapshot

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| application | application | |
| comment | string | Comment. Valid in POST. |
| components | array[components] | |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Creation time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | The snapshot name. Valid in POST. |
| svm | svm | |
| uuid | string | The snapshot UUID. Valid in URL. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create an application snapshot

POST `/application/applications/{application.uuid}/snapshots`

**Introduced In:** 9.6

Creates a snapshot of the application.

#### Required properties

- `name`

#### Recommended optional properties

- `consistency_type` - Track whether this snapshot is *application* or *crash* consistent.

#### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- DOC /application

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| application.uuid | string | path | True | Application UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| application | application | |
| comment | string | Comment. Valid in POST. |
| components | array[components] | |

| Name | Type | Description |
|---|---|---|
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Creation time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | The snapshot name. Valid in POST. |
| svm | svm | |
| uuid | string | The snapshot UUID. Valid in URL. |

**Example request**

```
{
  "application": {
    "name": "string",
    "uuid": "string"
  },
  "comment": "string",
  "components": [
    {
      "name": "string",
      "uuid": "string"
    }
  ],
  "consistency_type": "string",
  "create_time": "string",
  "name": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default, Error
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

application

| Name | Type | Description |
|---|---|---|
| name | string | Application name |
| uuid | string | The application UUID. Valid in URL. |

components

| Name | Type | Description |
|---|---|---|
| name | string | Component name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM name |
| uuid | string | SVM UUID |

application_snapshot

| Name | Type | Description |
|---|---|---|
| application | application | |
| comment | string | Comment. Valid in POST. |
| components | array[components] | |

| Name | Type | Description |
| --- | --- | --- |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Creation time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | The snapshot name. Valid in POST. |
| svm | svm | |
| uuid | string | The snapshot UUID. Valid in URL. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an application snapshot

DELETE `/application/applications/{application.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Delete a snapshot of an application

#### Query examples

Individual snapshots can be destroyed with no query parameters, or a range of snapshots can be destroyed at one time using a query.

The following query deletes all application snapshots created before May 4, 2017

```
DELETE
/application/applications/{application.uuid}/snapshots?create_time=<2017-
05-04T00:00:00-05:00
```

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an application snapshot

GET `/application/applications/{application.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Retrieve a snapshot of an application component.

This endpoint is only supported for Maxdata template applications.

Component snapshots are essentially more granular application snapshots. There is no difference beyond the scope of the operation.

#### Learn more

- DOC /application/applications/{application.uuid}/snapshots
- GET /application/applications/{uuid}/snapshots
- DOC /application

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | path | True | Snapshot UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | application | |
| comment | string | Comment. Valid in POST. |
| components | array[components] | |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |

| Name | Type | Description |
| --- | --- | --- |
| create_time | string | Creation time |
| is_partial | boolean | A partial snapshot means that not all volumes in an application component were included in the snapshot. |
| name | string | The snapshot name. Valid in POST. |
| svm | svm | |
| uuid | string | The snapshot UUID. Valid in URL. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "string"
  },
  "comment": "string",
  "components": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "string"
    }
  ],
  "consistency_type": "string",
  "create_time": "string",
  "name": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Application name |
| uuid | string | The application UUID. Valid in URL. |

components

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Component name |
| uuid | string | Component UUID |

svm

| Name | Type | Description |
|---|---|---|
| name | string | SVM name |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Restore an application snapshot**

POST `/application/applications/{application.uuid}/snapshots/{uuid}/restore`

**Introduced In:** 9.6

Restore an application snapshot

Restoring an application snapshot reverts all storage elements in the snapshot to the state in which the snapshot was in when the snapshot was taken. This restoration does not apply to access settings that might have changed since the snapshot was created.

**Learn more**

- DOC /application
- Asynchronous operations

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.uuid | string | path | True | Application UUID |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default, Error
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Delete an application and all associated data

DELETE `/application/applications/{uuid}`

**Introduced In:** 9.6

Deletes an application and all associated data.

### Warning - this deletes it all, including your data

This deletes everything created with the application, including any volumes, LUNs, NFS export policies, CIFS shares, and initiator groups. Initiator groups are only destroyed if they were created as part of an application and are no longer in use by other applications.

### Learn more

- DOC /application
- Asynchronous operations

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Application UUID |
| delete_data | boolean | query | False | By default, deleting an application deletes all of the application's data. By setting this parameter to "false", the application's data is preserved, but can no longer be managed through application APIs.<br>• Introduced in: 9.8<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

### Response

```
Status: 202, Accepted
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

#### See Definitions

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve an application

GET `/application/applications/{uuid}`

**Introduced In:** 9.6

Retrieves an application

### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `<template>` the property corresponding to the `template.name` of the application

### Property overview

An application includes three main groups or properties.

- Generic properties - such as the `name`, `template.name`, and `state` of the application. These properties are all inexpensive to retrieve and their meaning is consistent for every type of application.
- `statistics.*` - application statistics report live usage data about the application and its components. Various space and IOPS details are included at both the application level and at a per component level. The application model includes a detailed description of each property. These properties are slightly more expensive than the generic properties because live data must be collected from every storage element in the application.
- `<template>` - the property corresponding to the value of the `template.name` returns the contents of the application in the same layout that was used to provision the application. This information is very expensive to retrieve because it requires collecting information about all the storage and access settings for every element of the application. There are a few notable limitations to what can be returned in the `<template>` section:
  - The `new_igroups` array of many SAN templates is not returned by GET. This property allows igroup creation in the same call that creates an application, but is not a property of the application itself. The `new_igroups` array is allowed during PATCH operations, but that does not modify the `new_igroups` of the application. It is another way to allow igroup creation while updating the application to use a different igroup.
  - The `vdi_on_san` and `vdi_on_nas` `desktops.count` property is rounded to the nearest 1000 during creation, and is reported with that rounding applied.
  - The `mongo_db_on_san` `dataset.element_count` property is rounded up to an even number, and is reported with that rounding applied.
  - The `sql_on_san` and `sql_on_smb` `server_cores_count` property is limited to 8 for GET operations. Higher values are accepted by POST, but the impact of the `server_cores_count` property on the application layout currently reaches its limit at 8.

### Learn more

- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Application UUID |
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Application Name. This field is user supplied when the application is created. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |

| Name | Type | Description |
| --- | --- | --- |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| smart_container | boolean | Identifies if this is a smart container or not. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| svm | svm | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |

| Name | Type | Description |
| --- | --- | --- |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    },
    "snapshots": {
      "href": "/api/resourcelink"
    }
  },
  "creation_timestamp": "string",
  "delete_data": true,
  "generation": 0,
  "name": "string",
  "nas": {
    "application_components": [
      {}
    ],
    "cifs_access": [
      {
        "access": "string",
        "user_or_group": "string"
      }
    ],
    "cifs_share_name": "string",
    "exclude_aggregates": [
      {
        "name": "string",
        "uuid": "string"
      }
    ],
    "nfs_access": [
      {
        "access": "string",
        "host": "string"
      }
    ],
    "protection_type": {
      "local_policy": "string",
      "local_rpo": "string",
      "remote_rpo": "string"
    }
  },
  "nvme": {
    "components": [
```

```
      {
        "name": "string",
        "os_type": "string",
        "performance": {
          "storage_service": {
            "name": "string"
          }
        },
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "subsystem": {
          "hosts": [
            {
              "dh_hmac_chap": {
                "group_size": "string",
                "hash_function": "string"
              },
              "nqn": "string",
              "priority": "string"
            }
          ],
          "name": "string",
          "os_type": "string",
          "uuid": "string"
        },
        "tiering": {
          "policy": "string"
        }
      }
    ],
    "os_type": "string",
    "rpo": {
      "local": {
        "name": "string",
        "policy": "string"
      },
      "remote": {
        "name": "string"
      }
    }
  },
  "protection_granularity": "string",
```

```
  "rpo": {
    "components": [
      {
        "name": "string",
        "rpo": {
          "local": {
            "description": "string",
            "name": "string"
          },
          "remote": {
            "description": "string",
            "name": "string"
          }
        },
        "uuid": "string"
      }
    ],
    "local": {
      "description": "string",
      "name": "string"
    },
    "remote": {
      "description": "string",
      "name": "string"
    }
  },
  "s3_bucket": {
    "application_components": [
      {
        "access_policies": [
          {
            "actions": [
              "string"
            ],
            "conditions": [
              {
                "delimiters": [
                  "string"
                ],
                "max_keys": [
                  "integer"
                ],
                "operator": "string",
                "prefixes": [
                  "string"
                ],
```

```
                "source_ips": [
                  "string"
                ],
                "usernames": [
                  "string"
                ]
              }
            ],
            "effect": "string",
            "principals": [
              "string"
            ],
            "resources": [
              "string"
            ],
            "sid": "string"
          }
        ],
        "bucket_endpoint_type": "string",
        "comment": "string",
        "default_retention_period": "string",
        "exclude_aggregates": [
          {
            "name": "string",
            "uuid": "string"
          }
        ],
        "name": "string",
        "nas_path": "string",
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "retention_mode": "string",
        "snapshot_policy": {
          "name": "string",
          "uuid": "string"
        },
        "storage_service": {
          "name": "string"
        },
        "uuid": "string",
        "versioning_state": "string"
      }
```

```
      ],
      "protection_type": {
        "remote_rpo": "string"
      }
    },
    "san": {
      "application_components": [
        {
          "igroup_name": "string",
          "name": "string",
          "os_type": "string",
          "qos": {
            "policy": {
              "name": "string",
              "uuid": "string"
            }
          },
          "storage_service": {
            "name": "string"
          },
          "tiering": {
            "policy": "string"
          }
        }
      ],
      "exclude_aggregates": [
        {
          "name": "string",
          "uuid": "string"
        }
      ],
      "os_type": "string",
      "protection_type": {
        "local_policy": "string",
        "local_rpo": "string",
        "remote_rpo": "string"
      }
    },
    "state": "string",
    "statistics": {
      "components": [
        {
          "iops": {
            "per_tb": 0,
            "total": 0
          },
```

```
        "latency": {
          "average": 0,
          "raw": 0
        },
        "name": "string",
        "snapshot": {
          "reserve": 0,
          "used": 0
        },
        "space": {
          "available": 0,
          "logical_used": 0,
          "provisioned": 0,
          "reserved_unused": 0,
          "savings": 0,
          "used": 0,
          "used_excluding_reserves": 0,
          "used_percent": 0
        },
        "storage_service": {
          "name": "string",
          "uuid": "string"
        },
        "uuid": "string"
      }
    ],
    "iops": {
      "per_tb": 0,
      "total": 0
    },
    "latency": {
      "average": 0,
      "raw": 0
    },
    "snapshot": {
      "reserve": 0,
      "used": 0
    },
    "space": {
      "available": 0,
      "logical_used": 0,
      "provisioned": 0,
      "reserved_unused": 0,
      "savings": 0,
      "used": 0,
      "used_excluding_reserves": 0,
```

```
      "used_percent": 0
    }
  },
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "template": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "protocol": "string",
    "version": 0
  },
  "uuid": "string"
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |
| snapshots | href | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| element_count | integer | The number of storage elements (LUNs for SAN) of the database to maintain. Must be an even number between 2 and 16. Odd numbers will be rounded up to the next even number within range. |
| replication_factor | integer | The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

initiator_objects

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

secondary_igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
| --- | --- | --- |
| dataset | dataset | |
| os_type | string | The name of the host OS running the application. |
| primary_igroup_name | string | The initiator group for the primary. |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ID of an existing NFS export policy. |
| name | string | The name of an existing NFS export policy. |

component

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source component. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source SVM. |

origin

| Name | Type | Description |
| --- | --- | --- |
| component | component | |
| svm | svm | |

flexcache

| Name | Type | Description |
| --- | --- | --- |
| dr_cache | boolean | Dr-cache is a FlexCache volume create time option that has the same flexgroup-msid as that of the origin of a FlexCache volume. By default, dr-cache is disabled. The flexgroup-msid of the FlexCache volume does not need to be same as that of the origin of a FlexCache volume. |
| origin | origin | |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an existing QoS policy. |
| uuid | string | The UUID of an existing QoS policy. Usage: <UUID> |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as quot;Plt;num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period. |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| retention | retention | |
| snaplock_type | string | The SnapLock type of the smart container. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

nas_application_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
|---|---|---|
| export_policy | export_policy | |
| flexcache | flexcache | |
| name | string | The name of the application component. |
| qos | qos | |
| scale_out | boolean | Denotes a Flexgroup. |
| share_count | integer | The number of shares in the application component. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Indicates whether Snapshot copy locking is enabled on the volume. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

| Name | Type | Description |
|---|---|---|
| access | string | The CIFS access granted to the user or group. |
| user_or_group | string | The name of the CIFS user or group that will be granted access. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the aggregate to exclude. Usage: <aggregate name> |
| uuid | string | The ID of the aggregate to exclude. Usage: <UUID> |

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The NFS access granted. |
| host | string | The name of the NFS entity granted access. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

nas

A generic NAS application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| cifs_access | array[app_cifs_access] | The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access. |

| Name | Type | Description |
| --- | --- | --- |
| cifs_share_name | string | The name of the CIFS share. Usage: <Share> |
| exclude_aggregates | array[exclude_aggregates] | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

performance

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| qos | qos | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |

| Name | Type | Description |
|---|---|---|
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
|---|---|---|
| name | string | The local RPO of the application. |
| policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |

remote

| Name | Type | Description |
|---|---|---|
| name | string | The remote RPO of the application. |

rpo

| Name | Type | Description |
|---|---|---|
| local | local | |
| remote | remote | |

zapp_nvme

An NVME application.

| Name | Type | Description |
|---|---|---|
| components | array[components] | |
| os_type | string | The name of the host OS running the application. |
| rpo | rpo | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| mirrored | boolean | Specifies whether the redo log group should be mirrored. |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| ora_home | ora_home | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| copies | integer | The number of CRS volumes. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
|---|---|---|
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

local

| Name | Type | Description |
|---|---|---|
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the component. This indicates how often component Snapshot copies are automatically created. |

remote

| Name | Type | Description |
|---|---|---|
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | Component Name. |
| rpo | rpo | |
| uuid | string | Component UUID. |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the application. This indicates how often application Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the application. A remote RPO of zero indicates that the application is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| is_supported | boolean | Is RPO supported for this application? Generation 1 applications did not support Snapshot copies or MetroCluster. |
| local | local | |
| remote | remote | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| operator | string | Policy Condition Operator. |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[zapp_s3_bucket_application_components_access_policies_conditions] | conditions. |
| effect | string | Allow or Deny Access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Statement Identifier Usage: <(size 1..256)> |

snapshot_policy

| Name | Type | Description |
|---|---|---|
| name | string | The name of the snapshot policy that is used for the S3 bucket. |
| uuid | string | The UUID of the snapshot policy that is used for the S3 bucket. Usage: <UUID> |

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
| --- | --- | --- |
| access_policies | array[zapp_s3_bucket_application_components_access_policies] | The list of S3 objectstore policies to be created. |
| bucket_endpoint_type | string | The type of bucket. |
| capacity_tier | boolean | Prefer lower latency storage under similar media costs. |
| comment | string | Object Store Server Bucket Description Usage: <(size 1..256)> |
| default_retention_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as quot;Plt;num>Y" and quot;Plt;num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported. Usage: {{<integer> days\|years} \| none} |
| exclude_aggregates | array[exclude_aggregates] | |
| name | string | The name of the application component. |
| nas_path | string | The path to which the bucket corresponds to. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| retention_mode | string | The lock mode of the bucket. <br>compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects. <br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. For s3 type buckets, the default value is no_lock. |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| snapshot_policy | snapshot_policy | |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |
| versioning_state | string | Bucket Versioning State. For nas type buckets, this field is not set. For s3 type buckets, the default value is disabled. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| remote_rpo | string | The remote RPO of the application. |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |
| protection_type | protection_type | |

san_application_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
|---|---|---|
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| qos | qos | |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

san

A generic SAN application.

| Name | Type | Description |
|---|---|---|
| application_components | array[application_components] | |
| exclude_aggregates | array[exclude_aggregates] | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

access

| Name | Type | Description |
| --- | --- | --- |
| installer | string | SQL installer admin user name. |
| service_account | string | SQL service account user name. |

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| access | access | |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application component. |
| total | integer | The total number of IOPS being used by the application component. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this component. |
| raw | integer | The cumulative response time in microseconds for this component. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| reserve | integer | The amount of space reserved by the system for Snapshot copies. |
| used | integer | The amount of spacing currently in use by the system to store Snapshot copies. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application component. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application component. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service name. AFF systems support the extreme storage service. All other systems only support value. |
| uuid | string | The storage service UUID. |

components

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| latency | latency | |
| name | string | Component Name. |
| shared_storage_pool | boolean | An application component is considered to use a shared storage pool if storage elements for for other components reside on the same aggregate as storage elements for this component. |
| snapshot | snapshot | |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| statistics_incomplete | boolean | If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available. |
| storage_service | storage_service | |
| uuid | string | Component UUID. |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application. |
| total | integer | The total number of IOPS being used by the application. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this application. |
| raw | integer | The cumulative response time in microseconds for this application. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |
| used | integer | The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |

| Name | Type | Description |
| --- | --- | --- |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application. |

statistics

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| iops | iops | |
| latency | latency | |
| shared_storage_pool | boolean | An application is considered to use a shared storage pool if storage elements for multiple components reside on the same aggregate. |
| snapshot | snapshot | |
| space | space | |
| statistics_incomplete | boolean | If not all storage elements of the application are currently available, the returned statistics might only include data from those elements that were available. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | SVM Name. Either the SVM name or UUID must be provided to create an application. |
| uuid | string | SVM UUID. Either the SVM name or UUID must be provided to create an application. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

template

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The name of the template that was used to provision this application. |
| protocol | string | The protocol access of the template that was used to provision this application. |
| version | integer | The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| size | integer | The size of the desktops. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

hyper_v_access

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | Hyper-V service account. |

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of datastores to support. |
| size | integer | The size of the datastore. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| hypervisor | string | The name of the hypervisor hosting the application. |

| Name | Type | Description |
|---|---|---|
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Update application properties

PATCH `/application/applications/{uuid}`

**Introduced In:** 9.6

Updates the properties of an application.

### Overview

Similar to creating an application, modification is done using the template properties of an application. The `storage_service`, `size`, and `igroup_name` of an application may be modified.

`storage_service`

Storage service modifications are processed in place, meaning that the storage can not be moved to a location with more performance headroom to accommodate the request. If the current backing storage of the application is in a location that can support increased performance, the QoS policies associated with the application will be modified to allow it. If not, an error will be returned. A storage service modification to a lower tier of performance is always allowed, but the reverse modification may not be supported if the cluster is over provisioned and the cluster is unlikely to be able to fulfil the original storage service.

`size`

Size modifications are processed in a variety of ways depending on the type of application. For NAS applications, volumes are grown or new volumes are added. For SAN applications, LUNs are grown, new LUNs are added to existing volumes, or new LUNs are added to new volumes. If new storage elements are created, they can be found using the GET /application/applications/{application.uuid}/components interface. The creation time of each storage object is included, and the newly created objects will use the same naming scheme as the previous objects. Resize follows the best practices associated with the type of application being expanded. Reducing the size of an application is not supported.

`igroup_name`

Modification of the igroup name allows an entire application to be mapped from one initiator group to another. Data access will be interrupted as the LUNs are unmapped from the original igroup and remapped to the new one.

### Application state

During a modification, the `state` property of the application updates to indicate `modifying`. In `modifying` state, statistics are not available and snapshot operations are not allowed. If the modification fails, it is possible for the application to be left in an inconsistent state, with the underlying ONTAP storage elements not matching across a component. When this occurs, the application is left in the `modifying` state until the command is either retried and succeeds or a call to restore the original state is successful.

### Examples

1. Change the storage service of the database of the Oracle application to *extreme* and resize the redo logs to *100GB*.

```
  {
 "oracle_on_nfs": {
   "db": {
     "storage_service": {
       "name": "extreme"
     }
   },
   "redo_log": {
     "size": "100GB"
   }
 }
  }
```

2. Change the storage service, size, and igroup of a generic application by component name.

```
  {
"san": {
  "application_components": [
    {
      "name": "component1",
      "storage_service": {
        "name": "value"
      }
    },
    {
      "name": "component2",
      "size": "200GB"
    },
    {
      "name": "component3",
      "igroup_name": "igroup5"
    }
  ]
}
 }
```

### Learn more

- DOC /application
- Asynchronous operations

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Application UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |

| Name | Type | Description |
| --- | --- | --- |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Response

```
Status: 202, Accepted
```

### Error

```
Status: Default, Error
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of an igroup to nest within a parent igroup. Mutually exclusive with initiators and initiator_objects. |
| uuid | string | The UUID of an igroup to nest within a parent igroup Usage: <UUID> |

initiator_objects

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| name | string | The WWPN, IQN, or Alias of the initiator. Mutually exclusive with nested igroups and the initiators array. |

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

protection_type

| Name | Type | Description |
|---|---|---|
| local_rpo | string | The local RPO of the application. |

secondary_igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
|---|---|---|
| dataset | dataset | |
| new_igroups | array[mongo_db_on_san_new_igroups] | The list of initiator groups to create. |
| primary_igroup_name | string | The initiator group for the primary. |

| Name | Type | Description |
| --- | --- | --- |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

component

svm

origin

flexcache

policy

qos

retention

snaplock

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

nas_application_components_tiering

application-components.tiering

application_components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

exclude_aggregates

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

nas

A generic NAS application.

| Name | Type | Description |
|---|---|---|
| application_components | array[application_components] | |
| protection_type | protection_type | |

performance

| Name | Type | Description |
|---|---|---|
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
|---|---|---|
| controller_secret_key | string | Authentication Controller Secret. |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |
| host_secret_key | string | Authentication Host Secret. |

hosts

| Name | Type | Description |
|---|---|---|
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
|---|---|---|
| name | string | The local RPO of the application. |

remote

rpo

| Name | Type | Description |
|---|---|---|
| local | local | |

zapp_nvme

An NVME application.

| Name | Type | Description |
|---|---|---|
| components | array[components] | |
| rpo | rpo | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[oracle_on_san_new_igroups] | The list of initiator groups to create. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| new_igroups | array[oracle_rac_on_san_new_igroups] | The list of initiator groups to create. |

| Name | Type | Description |
| --- | --- | --- |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the component. This indicates how often component Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the component. A remote RPO of zero indicates that the component is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| local | local | |
| remote | remote | |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | Component Name. |
| rpo | rpo | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Component UUID. |

local

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the local RPO. This will include details about the Snapshot copy schedule. |
| name | string | The local RPO of the application. This indicates how often application Snapshot copies are automatically created. |

remote

| Name | Type | Description |
| --- | --- | --- |
| description | string | A detailed description of the remote RPO. |
| name | string | The remote RPO of the application. A remote RPO of zero indicates that the application is synchronously replicated to another cluster. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| is_supported | boolean | Is RPO supported for this application? Generation 1 applications did not support Snapshot copies or MetroCluster. |
| local | local | |
| remote | remote | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| principals | array[string] | |
| resources | array[string] | |

snapshot_policy

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the application component. |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
|---|---|---|
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |

san_application_components_tiering

application-components.tiering

application_components

| Name | Type | Description |
|---|---|---|
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

san

A generic SAN application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| new_igroups | array[san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| os_type | string | The name of the host OS accessing the application. The default value is the host OS that is running the application. |
| protocol | string | The protocol of the new initiator group. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| new_igroups | array[sql_on_san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |
| temp_db | temp_db | |

access

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| temp_db | temp_db | |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application component. |
| total | integer | The total number of IOPS being used by the application component. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this component. |
| raw | integer | The cumulative response time in microseconds for this component. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| reserve | integer | The amount of space reserved by the system for Snapshot copies. |
| used | integer | The amount of spacing currently in use by the system to store Snapshot copies. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application component. |

| Name | Type | Description |
| --- | --- | --- |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |
| used | integer | The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application component. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service name. AFF systems support the extreme storage service. All other systems only support value. |
| uuid | string | The storage service UUID. |

components

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| latency | latency | |
| name | string | Component Name. |

| Name | Type | Description |
| --- | --- | --- |
| shared_storage_pool | boolean | An application component is considered to use a shared storage pool if storage elements for for other components reside on the same aggregate as storage elements for this component. |
| snapshot | snapshot | |
| space | space | |
| statistics_incomplete | boolean | If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available. |
| storage_service | storage_service | |
| uuid | string | Component UUID. |

iops

| Name | Type | Description |
| --- | --- | --- |
| per_tb | integer | The number of IOPS per terabyte of logical space currently being used by the application. |
| total | integer | The total number of IOPS being used by the application. |

latency

| Name | Type | Description |
| --- | --- | --- |
| average | integer | The cumulative average response time in microseconds for this application. |
| raw | integer | The cumulative response time in microseconds for this application. |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| logical_used | integer | The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used. |
| provisioned | integer | The originally requested amount of space that was provisioned for the application. |
| reserved_unused | integer | The amount of space reserved for system features such as Snapshot copies that has not yet been used. |
| savings | integer | The amount of space saved by all enabled space saving features. |
| used | integer | The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies. |
| used_excluding_reserves | integer | The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies. |

| Name | Type | Description |
| --- | --- | --- |
| used_percent | integer | The percentage of the originally provisioned space that is currently being used by the application. |

statistics

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| iops | iops | |
| latency | latency | |
| shared_storage_pool | boolean | An application is considered to use a shared storage pool if storage elements for multiple components reside on the same aggregate. |
| snapshot | snapshot | |
| space | space | |
| statistics_incomplete | boolean | If not all storage elements of the application are currently available, the returned statistics might only include data from those elements that were available. |

self_link

template

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocol access of the template that was used to provision this application. |

| Name | Type | Description |
| --- | --- | --- |
| version | integer | The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| storage_service | storage_service | |

hyper_v_access

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| protocol | string | The protocol of the new initiator group. |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[vdi_on_san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of datastores to support. |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroups] | |
| initiator_objects | array[initiator_objects] | |
| initiators | array[string] | |
| name | string | The name of the new initiator group. |
| protocol | string | The protocol of the new initiator group. |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
| --- | --- | --- |
| datastore | datastore | |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| new_igroups | array[vsi_on_san_new_igroups] | The list of initiator groups to create. |
| protection_type | protection_type | |

application

Applications

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| creation_timestamp | string | The time when the application was created. |
| generation | integer | The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |

| Name | Type | Description |
| --- | --- | --- |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protection_granularity | string | Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components. |
| rpo | rpo | |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| state | string | The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported. |
| statistics | statistics | |
| template | template | |
| uuid | string | Application UUID. This field is generated when the application is created. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage application consistency groups

### Manage application consistency groups

#### Overview

A consistency group is a group of volumes that supports capabilities such as creating a snapshot of all of its member volumes at the same point-in-time with a write-fence, thus ensuring a consistent image of the volumes at that time.

Applications with datasets scoped to a single volume can have its contents saved to a snapshot, replicated, or cloned in a crash-consistent manner implicitly with corresponding native ONTAP volume-granular operations. Applications with datasets spanning a group of multiple volumes must have such operations performed on the group. Typically, by first fencing writes to all the volumes in the group, flushing any writes pending in queues, executing the intended operation, that is, take snapshot of every volume in the group and when that is complete, unfence and resume writes. A consistency group is the conventional mechanism for providing such group semantics.

#### Consistency group APIs

The following APIs are used to perform operations related to consistency groups:

– GET /api/application/consistency-groups

– POST /api/application/consistency-groups

– GET /api/application/consistency-groups/{uuid}

– PATCH /api/application/consistency-groups/{uuid}

– DELETE /api/application/consistency-groups/{uuid}

#### Examples

### Retrieving all consistency groups of an SVM

```
# The API:
/api/application/consistency-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/application/consistency-
groups?svm.name=vs1" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/6f48d798-0a7f-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
    "name": "parent_cg",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b22c85-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
```

```
    "uuid": "c1b270b1-0a82-11ec-a449-005056bbcf9f",
    "name": "child_1",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b270b1-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b270c3-0a82-11ec-a449-005056bbcf9f",
    "name": "child_2",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b270c3-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/application/consistency-groups"
  }
}
}
```

**Retrieving details of all consistency groups of an SVM**

Retrieving details of the consistency groups for a specified SVM. These details are considered to be performant and will return within 1 second when 40 records or less are requested.

```
curl -X GET "https://<mgmt-ip>/api/application/consistency-
groups?svm.name=vs1&fields=*&max_records=40"

#### Response:
{
"records": [
  {
    "uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
    "name": "vol1",
    "svm": {
      "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
      "name": "vs1",
      "_links": {
```

```
        "self": {
          "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
        }
      }
    },
    "space": {
      "size": 108003328,
      "available": 107704320,
      "used": 299008
    },
    "replicated": false,
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/6f48d798-0a7f-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
    "name": "parent_cg",
    "svm": {
      "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
        }
      }
    },
    "snapshot_policy": {
      "name": "default-1weekly",
      "uuid": "a30bd0fe-067d-11ec-a449-005056bbcf9f",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/a30bd0fe-067d-11ec-a449-
005056bbcf9f"
        }
      }
    },
    "consistency_groups": [
      {
        "uuid": "c1b270b1-0a82-11ec-a449-005056bbcf9f",
        "name": "child_1",
        "space": {
          "size": 41943040,
```

```
          "available": 39346176,
          "used": 499712
        },
        "_links": {
          "self": {
            "href": "/api/application/consistency-groups/c1b270b1-0a82-
11ec-a449-005056bbcf9f"
          }
        }
      },
      {
        "uuid": "c1b270c3-0a82-11ec-a449-005056bbcf9f",
        "name": "child_2",
        "space": {
          "size": 41943040,
          "available": 39350272,
          "used": 495616
        },
        "_links": {
          "self": {
            "href": "/api/application/consistency-groups/c1b270c3-0a82-
11ec-a449-005056bbcf9f"
          }
        }
      }
    ],
    "space": {
      "size": 83886080,
      "available": 78696448,
      "used": 995328
    },
    "replicated": false,
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b22c85-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b270b1-0a82-11ec-a449-005056bbcf9f",
    "name": "child_1",
    "parent_consistency_group": {
      "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
      "name": "parent_cg",
      "_links": {
```

```
        "self": {
          "href": "/api/application/consistency-groups/c1b22c85-0a82-11ec-
a449-005056bbcf9f"
        }
      }
    },
    "svm": {
      "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
        }
      }
    },
    "snapshot_policy": {
      "name": "default",
      "uuid": "a30b60a4-067d-11ec-a449-005056bbcf9f",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/a30b60a4-067d-11ec-a449-
005056bbcf9f"
        }
      }
    },
    "space": {
      "size": 41943040,
      "available": 39346176,
      "used": 499712
    },
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b270b1-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b270c3-0a82-11ec-a449-005056bbcf9f",
    "name": "child_2",
    "parent_consistency_group": {
      "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
      "name": "parent_cg",
      "_links": {
        "self": {
          "href": "/api/application/consistency-groups/c1b22c85-0a82-11ec-
```

```
a449-005056bbcf9f"
        }
      }
    },
    "svm": {
      "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
        }
      }
    },
    "snapshot_policy": {
      "name": "default",
      "uuid": "a30b60a4-067d-11ec-a449-005056bbcf9f",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/a30b60a4-067d-11ec-a449-
005056bbcf9f"
        }
      }
    },
    "space": {
      "size": 41943040,
      "available": 39350272,
      "used": 495616
    },
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b270c3-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/application/consistency-
groups?svm.name=vs1&fields=*&max_records=40"
  }
}
}
```

#### Retrieving details of non-nested consistency groups

Retrieves details of the consistency groups without nested consistency groups, or only the parent consistency group for a number of consistency groups of a specified SVM.

```
curl -X GET "https://<mgmt-ip>/api/application/consistency-
groups?svm.name=vs1&parent_consistency_group.uuid=null"

#### Response:
{
"records": [
  {
    "uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
    "name": "vol1",
    "svm": {
      "name": "vs1"
    },
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/6f48d798-0a7f-11ec-
a449-005056bbcf9f"
      }
    }
  },
  {
    "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
    "name": "parent_cg",
    "svm": {
      "name": "vs1"
    },
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/c1b22c85-0a82-11ec-
a449-005056bbcf9f"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/application/consistency-
groups?svm.name=vs1&parent_consistency_group.uuid=null"
  }
}
}
```

#### Creating a single consistency group with a new SAN volume

Provisions an application with one consistency group, each with one new SAN volumes, with one LUN, an igroup and no explicit snapshot policy, FabricPool tiering policy, storage service, and QoS policy specification. The igroup to map a LUN to is specified at LUN-granularity.

```
curl -X POST https://<mgmt-ip>/api/application/consistency-
groups?return_records=true -d '{ "svm": { "name": "vs1" }, "luns": [ {
"name": "/vol/vol1/lun1", "space": { "size": "100mb" }, "os_type":
"linux", "lun_maps": [ { "igroup": { "name": "igroup1", "initiators": [ {
"name": "example_name" } ] } } ] } ] }'

#### Response:
{
"num_records": 1,
"records": [
  {
    "uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
    "name": "vol1",
    "svm": {
    "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
      }
    }
    },
    "luns": [
    {
      "lun_maps": [
        {
          "igroup": {
            "name": "igroup1",
            "initiators": [
              {
                "name": "example_name"
              }
            ]
          }
        }
      ],
      "name": "/vol/vol1/lun1",
      "os_type": "linux",
      "space": {
        "size": 104857600
      }
```

```
    }
    ]
  }
],
"job": {
  "uuid": "6f4907ae-0a7f-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
    "href": "/api/cluster/jobs/6f4907ae-0a7f-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Creating an Application with two consistency groups with existing SAN volumes**

Provisions an application with two consistency groups, each with two existing SAN volumes, a snapshot policy at application-granularity, and a distinct consistency group granular snapshot policy.

```
curl -X POST https://<mgmt-ip>/api/application/consistency-
groups?return_records=true -d '{ "svm": { "name": "vs1" }, "name":
"parent_cg", "snapshot_policy": { "name": "default-1weekly" },
"consistency_groups": [ { "name": "child_1", "snapshot_policy": { "name":
"default" }, "volumes": [ { "name": "existing_vol1",
"provisioning_options": { "action": "add" } }, { "name": "existing_vol2",
"provisioning_options": { "action": "add" } } ] }, { "name": "child_2",
"snapshot_policy": { "name": "default" }, "volumes": [ { "name":
"existing_vol3", "provisioning_options": { "action": "add" } }, { "name":
"existing_vol4", "provisioning_options": { "action": "add" } } ] } ] }'

#### Response:
{
"num_records": 1,
"records": [
  {
    "uuid": "c1b22c85-0a82-11ec-a449-005056bbcf9f",
    "name": "parent_cg",
    "svm": {
      "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
        }
      }
    },
```

```
    "snapshot_policy": {
      "name": "default-1weekly"
    },
    "consistency_groups": [
      {
        "uuid": "c1b270b1-0a82-11ec-a449-005056bbcf9f",
        "name": "child_1",
        "snapshot_policy": {
          "name": "default"
        },
        "volumes": [
          {
            "name": "existing_vol1"
          },
          {
            "name": "existing_vol2"
          }
        ]
      },
      {
        "uuid": "c1b270c3-0a82-11ec-a449-005056bbcf9f",
        "name": "child_2",
        "snapshot_policy": {
          "name": "default"
        },
        "volumes": [
          {
            "name": "existing_vol3"
          },
          {
            "name": "existing_vol4"
          }
        ]
      }
    ]
  }
],
"job": {
  "uuid": "c1b272b9-0a82-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c1b272b9-0a82-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Retrieving specific details of an existing consistency group**

Retrieves the details of an existing consistency group.

```
curl -X GET https://<mgmt-ip>/api/application/consistency-groups/6f48d798-
0a7f-11ec-a449-005056bbcf9f

#### Response:
{
"uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
"name": "vol1",
"svm": {
  "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
    }
  }
},
"space": {
  "size": 108003328,
  "available": 107724800,
  "used": 278528
},
"replicated": false,
"_links": {
  "self": {
    "href": "/api/application/consistency-groups/6f48d798-0a7f-11ec-a449-
005056bbcf9f"
  }
}
}
```

**Retrieving all details of an existing consistency group**

Retrieves all details of an existing consistency group. These details are expensive and are not guaranteed to return within one second.

```
curl -X GET https://<mgmt-ip>/api/application/consistency-groups/6f48d798-
0a7f-11ec-a449-005056bbcf9f?fields=**

#### Response:
{
"uuid": "6f48d798-0a7f-11ec-a449-005056bbcf9f",
"name": "vol1",
```

```
"svm": {
  "uuid": "4853f97a-0a63-11ec-a449-005056bbcf9f",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/4853f97a-0a63-11ec-a449-005056bbcf9f"
    }
  }
},
"qos": {
  "policy": {
    "uuid": "b7189398-e572-48ab-8f69-82cd46580812",
    "name": "extreme-fixed",
    "_links": {
      "self": {
        "href": "/api/storage/qos/policies/b7189398-e572-48ab-8f69-
82cd46580812"
      }
    }
  }
},
"tiering": {
  "policy": "none"
},
"create_time": "2021-08-31T13:18:24-04:00",
"volumes": [
  {
    "uuid": "6f516c6c-0a7f-11ec-a449-005056bbcf9f",
    "qos": {
      "policy": {
        "uuid": "b7189398-e572-48ab-8f69-82cd46580812",
        "name": "extreme-fixed",
        "_links": {
          "self": {
            "href": "/api/storage/qos/policies/b7189398-e572-48ab-8f69-
82cd46580812"
          }
        }
      }
    },
    "tiering": {
      "policy": "none"
    },
    "comment": "",
    "create_time": "2021-08-31T13:18:22-04:00",
    "name": "vol1",
```

```
    "snapshot_policy": {
      "name": "default",
      "uuid": "a30b60a4-067d-11ec-a449-005056bbcf9f"
    },
    "space": {
      "size": 108003328,
      "available": 107569152,
      "used": 434176,
      "snapshot": {
        "used": 151552,
        "reserve_percent": 0,
        "autodelete_enabled": false
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/6f516c6c-0a7f-11ec-a449-
005056bbcf9f"
      }
    }
  }
],
"luns": [
  {
    "uuid": "6f51748a-0a7f-11ec-a449-005056bbcf9f",
    "location": {
      "logical_unit": "lun1",
      "node": {
        "name": "example_node_name",
        "uuid": "6eb682f2-067d-11ec-a449-005056bbcf9f",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/6eb682f2-067d-11ec-a449-
005056bbcf9f"
          }
        }
      },
      "volume": {
        "uuid": "6f516c6c-0a7f-11ec-a449-005056bbcf9f",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/6f516c6c-0a7f-11ec-a449-
005056bbcf9f"
          }
        }
```

```
      }
    },
    "lun_maps": [
      {
        "igroup": {
          "uuid": "6f4a4b86-0a7f-11ec-a449-005056bbcf9f",
          "name": "igroup1",
          "os_type": "linux",
          "protocol": "mixed",
          "initiators": [
            {
              "name": "example_name"
            }
          ],
          "_links": {
            "self": {
              "href": "/api/protocols/san/igroups/6f4a4b86-0a7f-11ec-a449-
005056bbcf9f"
            }
          }
        },
        "logical_unit_number": 0
      }
    ],
    "name": "/vol/vol1/lun1",
    "auto_delete": false,
    "class": "regular",
    "create_time": "2021-08-31T13:18:24-04:00",
    "os_type": "linux",
    "serial_number": "wIqM6]RfQK3t",
    "space": {
      "size": 104857600,
      "used": 0,
      "guarantee": {
        "requested": false,
        "reserved": false
      }
    },
    "status": {
      "container_state": "online",
      "mapped": true,
      "read_only": false,
      "state": "online"
    },
    "_links": {
      "self": {
```

```
        "href": "/api/storage/luns/6f51748a-0a7f-11ec-a449-005056bbcf9f"
      }
    }
  }
],
"space": {
  "size": 108003328,
  "available": 107569152,
  "used": 434176
},
"replicated": false,
"_links": {
  "self": {
    "href": "/api/application/consistency-groups/6f48d798-0a7f-11ec-a449-
005056bbcf9f?fields=**"
  }
}
}
```


#### Adding LUNs to an existing volume in an existing consistency group

Adds two NVMe namespaces to an existing volume in an existing consistency group, creates a new subsystem, and binds the new namespaces to it.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f48d798-0a7f-11ec-a449-005056bbcf9f' -d '{ "luns": [ { "name":
"/vol/vol1/new_luns", "provisioning_options": { "count": 2, "action":
"create" }, "space": { "size": "100mb" }, "os_type": "linux", "lun_maps":
[ { "igroup": { "name": "igroup2", "initiators": [ { "name":
"01:02:03:04:05:06:07:01" } ] } } ] } ] }'

#### Response:
{
"job": {
  "uuid": "5306ea44-0a87-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/5306ea44-0a87-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Restoring a consistency group to the contents of an existing snapshot

Restores an existing consistency group to the contents of an existing snapshot of the consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "restore_to": {
"snapshot": { "uuid": "92c6c770-17a1-11eb-b141-005056acd498"} } }' -H
"Accept: Application/hal+json"

#### Response:
{
"job": {
  "uuid": "8907bd9e-1463-11eb-a719-005056ac70af",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8907bd9e-1463-11eb-a719-005056ac70af"
    }
  }
}
}
```

#### Deleting a consistency group

Deletes a consistency group, where all storage originally associated with that consistency group remains in place.

```
curl -X DELETE 'https://<mgmt-ip>/api/application/consistency-
groups/6f48d798-0a7f-11ec-a449-005056bbcf9f'

#### Response:
{
}
```

#### Cloning an existing consistency group

The following example clones an existing consistency group with the current contents:

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-groups' -d '{
"name": "clone01_of_cg01","svm": { "name": "vs_0"},"clone": { "volume": {
"prefix": "my_clone_pfx","suffix": "my_clone_sfx"},"split_initiated":
true,"parent_consistency_group": { "name": "cg01","uuid": "ca5e76fb-98c0-
11ec-855a-005056a7693b"},"guarantee": { "type": "none"} } }' -H "accept:
application/hal+json"

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Cloning a consistency group from an existing snapshot**

The following example clones an existing consistency group with contents from an existing snapshot:

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-groups' -d '{
"name": "clone01_of_cg01","svm": { "name": "vs_0"},"clone": { "volume": {
"prefix": "my_clone_pfx","suffix": "my_clone_sfx"},"split_initiated":
true,"parent_snapshot":{ "name":
"snap01_of_cg01"},"parent_consistency_group": { "name": "cg01","uuid":
"ca5e76fb-98c0-11ec-855a-005056a7693b"},"guarantee": { "type": "none"} }
}' -H "accept: application/hal+json"

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

### Adding namespaces to an existing volume in an existing consistency group

To add two NVMe Namespaces to an existing volume in an existing consistency group, create a new subsystem and bind the new namespaces to it.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "namespaces": [ {
"name": "/vol/vol1/new_namespace", "space": { "size": "10M" }, "os_type":
"windows", "provisioning_options": { "count": 2 }, "subsystem_map": {
"subsystem": { "name": "mySubsystem", "hosts": [ { "nqn": "nqn.1992-
08.com.netapp:sn.d04594ef915b4c73b642169e72e4c0b1:subsystem.host1" }, {
"nqn": "nqn.1992-
08.com.netapp:sn.d04594ef915b4c73b642169e72e4c0b1:subsystem.host2" } ] } }
} ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

### Add a new volume in an existing consistency group

The following example adds two new volumes to an existing consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "volumes" : [ { "name":
"new_vol_", "provisioning_options": { "count" : "2"}, "space": { "size":
"1gb"} } ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Adding an existing volume to an existing consistency group

The following example adds an existing volume to an existing consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "volumes" : [ { "name":
"existing_vol", "provisioning_options": { "action" : "add"} } ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Promoting a single consistency group to a nested consistency group

The following example promotes a single consistency group to a nested consistency group with a new child consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "provisioning_options":
{ "action":"promote" }, "consistency_groups": [ { "name":"cg_child",
"provisioning_options": { "action":"create" } } ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Demoting a nested consistency group to a single consistency group

The following example demotes (flattens) a nested consistency group to a single consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "provisioning_options":
{ "action":"demote" } }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Adding a new child consistency group to nested consistency group

The following example adds a new child consistency group to an existing nested consistency group, creating a new volume.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "consistency_groups": [
{ "name": "cg_child5", "provisioning_options": { "action" : "create" },
"volumes": [ { "name": "child5_vol_1", "space": { "size": "100mb" } } ] }
] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Removing a child consistency group from nested consistency group

The following example removes a child consistency group from a nested consistency, changing it to a single consistency group with a new consistency group name.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "consistency_groups": [
{ "name": "cg_child5", "provisioning_options": { "action" : "remove",
"name" : "new_single_cg" } } ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

#### Create a new parent consistency group with an existing consistency group

The following example creates a new nested consistency group with an existing consistency group as child consistency group.

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-groups' -d '{
"svm": { "name": "vs1" }, "name": "cg_parent2", "consistency_groups": [ {
"name": "cg_large", "provisioning_options":{ "action": "add" } }, {
"name": "cg_standalone2", "provisioning_options":{ "action": "add" } } ]
}'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Reassign a volume to another child consistency group.**

The following example reassigns a volume from a child consistency group to another child consistency group with the same parent consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/6f51748a-0a7f-11ec-a449-005056bbcf9f' -d '{ "consistency_groups": [
{ "name": "cg_child1", "volumes": [ { "name": "child2_vol_1",
"provisioning_options": { "action":"reassign"} } ] }, { "name":
"cg_child2" } ] }'

#### Response:
{
"job": {
  "uuid": "8c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Retrieve details of a collection or consistency group**

GET `/application/consistency-groups`

**Introduced In:** 9.10

Retrieve details of a collection or a specific consistency group.

#### Notes

When volume granular properties, such as, the storage SLC, Fabric Pool tiering are not the same for all the existing volumes of a consistency group, the corresponding property is not reported at consistency group granularity. It is only reported if all the volumes of the consistency group have the same value for that property.

If this consistency group instance has 1 or more replication relationships, the "replicated" parameter is true. If there are no associated replication relationships, it is false. This parameter is only included in the output for Single-CG and Parent-CG, not for Child-CG.
If this consistency group instance has 1 or more replication relationships, the "replication_relationships" parameter is included in the output for Single-CG and Parent-CG instances. If there are no associated replication relationships, this parameter is not included in the output.
Note that this parameter is an array and as such it has as many elements as the number of replication relationships associated with this consistency group. Each element of the array describes properties of one replication relationship associated with this consistency group. The "uuid" parameter identifies a specific replication relationship and the "href" parameter is a link to the corresponding SnapMirror relationship. The "is_source" parameter is true if this consistency group is the source in that relationship, otherwise it is false.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See DOC Requesting specific fields to learn more.

- `volumes`
- `luns`
- `namespaces`

#### Related ONTAP commands

- `vserver consistency-group show'

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| application.component_type | string | query | False | Filter by application.component_type<br>• Introduced in: 9.12 |
| application.type | string | query | False | Filter by application.type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| tiering.policy | string | query | False | Filter by tiering.policy |
| luns.create_time | string | query | False | Filter by luns.create_time |
| luns.comment | string | query | False | Filter by luns.comment<br>• maxLength: 254<br>• minLength: 0 |
| luns.os_type | string | query | False | Filter by luns.os_type |
| luns.uuid | string | query | False | Filter by luns.uuid |
| luns.space.used | integer | query | False | Filter by luns.space.used |
| luns.space.guarantee.requested | boolean | query | False | Filter by luns.space.guarantee.requested<br>• Introduced in: 9.11 |
| luns.space.guarantee.reserved | boolean | query | False | Filter by luns.space.guarantee.reserved<br>• Introduced in: 9.11 |
| luns.space.snapshot.reserve_available | integer | query | False | Filter by luns.space.snapshot.reserve_available<br>• Introduced in: 9.18 |
| luns.space.snapshot.reserve_percent | integer | query | False | Filter by luns.space.snapshot.reserve_percent<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.space.snapshot.space_used_percent | integer | query | False | Filter by luns.space.snapshot.space_used_percent<br>• Introduced in: 9.18 |
| luns.space.snapshot.used | integer | query | False | Filter by luns.space.snapshot.used<br>• Introduced in: 9.18 |
| luns.space.snapshot.reserve_size | integer | query | False | Filter by luns.space.snapshot.reserve_size<br>• Introduced in: 9.18 |
| luns.space.snapshot.autodelete.enabled | boolean | query | False | Filter by luns.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| luns.space.size | integer | query | False | Filter by luns.space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |
| luns.serial_number | string | query | False | Filter by luns.serial_number<br>• maxLength: 12<br>• minLength: 12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.qos.policy.min_throughput_iops | integer | query | False | Filter by luns.qos.policy.min_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| luns.qos.policy.uuid | string | query | False | Filter by luns.qos.policy.uuid |
| luns.qos.policy.max_throughput_mbps | integer | query | False | Filter by luns.qos.policy.max_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| luns.qos.policy.name | string | query | False | Filter by luns.qos.policy.name |
| luns.qos.policy.min_throughput_mbps | integer | query | False | Filter by luns.qos.policy.min_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| luns.qos.policy.max_throughput_iops | integer | query | False | Filter by luns.qos.policy.max_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| luns.name | string | query | False | Filter by luns.name |
| luns.enabled | boolean | query | False | Filter by luns.enabled |
| luns.lun_maps.logical_unit_number | integer | query | False | Filter by luns.lun_maps.logical_unit_number |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.lun_maps.igroup.igroups.uuid | string | query | False | Filter by luns.lun_maps.igroup.igroups.uuid |
| luns.lun_maps.igroup.igroups.name | string | query | False | Filter by luns.lun_maps.igroup.igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| luns.lun_maps.igroup.uuid | string | query | False | Filter by luns.lun_maps.igroup.uuid |
| luns.lun_maps.igroup.protocol | string | query | False | Filter by luns.lun_maps.igroup.protocol |
| luns.lun_maps.igroup.name | string | query | False | Filter by luns.lun_maps.igroup.name<br>• maxLength: 96<br>• minLength: 1 |
| luns.lun_maps.igroup.os_type | string | query | False | Filter by luns.lun_maps.igroup.os_type |
| luns.lun_maps.igroup.comment | string | query | False | Filter by luns.lun_maps.igroup.comment<br>• Introduced in: 9.11<br>• maxLength: 254<br>• minLength: 0 |
| luns.lun_maps.igroup.initiators.name | string | query | False | Filter by luns.lun_maps.igroup.initiators.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.lun_maps.igroup.initiators.comment | string | query | False | Filter by luns.lun_maps.igroup.initiators.comment<br>• maxLength: 254<br>• minLength: 0 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.13 |
| statistics.used_space | integer | query | False | Filter by statistics.used_space<br>• Introduced in: 9.13 |
| statistics.throughput_raw.other | integer | query | False | Filter by statistics.throughput_raw.other<br>• Introduced in: 9.13 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.13 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br>• Introduced in: 9.13 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.13 |
| statistics.available_space | integer | query | False | Filter by statistics.available_space<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.13 |
| statistics.size | integer | query | False | Filter by statistics.size<br><br>• Introduced in: 9.13 |
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid |
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name |
| replication_source | boolean | query | False | Filter by replication_source |
| name | string | query | False | Filter by name |
| replicated | boolean | query | False | Filter by replicated |
| clone.guarantee.type | string | query | False | Filter by clone.guarantee.type<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| clone.parent_snapshot.name | string | query | False | Filter by clone.parent_snapshot.name<br><br>• Introduced in: 9.12 |
| clone.parent_snapshot.uuid | string | query | False | Filter by clone.parent_snapshot.uuid<br><br>• Introduced in: 9.15 |
| clone.parent_svm.name | string | query | False | Filter by clone.parent_svm.name<br><br>• Introduced in: 9.15 |
| clone.parent_svm.uuid | string | query | False | Filter by clone.parent_svm.uuid<br><br>• Introduced in: 9.15 |
| clone.volume.suffix | string | query | False | Filter by clone.volume.suffix<br><br>• Introduced in: 9.12 |
| clone.volume.prefix | string | query | False | Filter by clone.volume.prefix<br><br>• Introduced in: 9.12 |
| clone.split_complete_percent | integer | query | False | Filter by clone.split_complete_percent<br><br>• Introduced in: 9.15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| clone.parent_consistency_group.name | string | query | False | Filter by clone.parent_consistency_group.name<br>• Introduced in: 9.12 |
| clone.parent_consistency_group.uuid | string | query | False | Filter by clone.parent_consistency_group.uuid<br>• Introduced in: 9.12 |
| clone.is_flexclone | boolean | query | False | Filter by clone.is_flexclone<br>• Introduced in: 9.15 |
| clone.split_initiated | boolean | query | False | Filter by clone.split_initiated<br>• Introduced in: 9.12 |
| clone.split_estimate | integer | query | False | Filter by clone.split_estimate<br>• Introduced in: 9.15 |
| metric.size | integer | query | False | Filter by metric.size<br>• Introduced in: 9.13 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.13 |
| metric.available_space | integer | query | False | Filter by metric.available_space<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.13 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.13 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.13 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.13 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.13 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.13 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.13 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.13 |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other<br><br>• Introduced in: 9.13 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.13 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.13 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.13 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.13 |
| metric.used_space | integer | query | False | Filter by metric.used_space<br><br>• Introduced in: 9.13 |
| consistency_groups.uuid | string | query | False | Filter by consistency_groups.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.space.used | integer | query | False | Filter by consistency_groups.space.used |
| consistency_groups.space.size | integer | query | False | Filter by consistency_groups.space.size |
| consistency_groups.space.available | integer | query | False | Filter by consistency_groups.space.available |
| consistency_groups.snapshot_policy.uuid | string | query | False | Filter by consistency_groups.snapshot_policy.uuid |
| consistency_groups.snapshot_policy.name | string | query | False | Filter by consistency_groups.snapshot_policy.name |
| consistency_groups.luns.create_time | string | query | False | Filter by consistency_groups.luns.create_time |
| consistency_groups.luns.comment | string | query | False | Filter by consistency_groups.luns.comment<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.luns.os_type | string | query | False | Filter by consistency_groups.luns.os_type |
| consistency_groups.luns.uuid | string | query | False | Filter by consistency_groups.luns.uuid |
| consistency_groups.luns.space.used | integer | query | False | Filter by consistency_groups.luns.space.used |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.space.guarantee.requested | boolean | query | False | Filter by consistency_groups.luns.space.guarantee.requested<br><br>• Introduced in: 9.11 |
| consistency_groups.luns.space.guarantee.reserved | boolean | query | False | Filter by consistency_groups.luns.space.guarantee.reserved<br><br>• Introduced in: 9.11 |
| consistency_groups.luns.space.snapshot.reserve_available | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_available<br><br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.reserve_percent | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.space_used_percent | integer | query | False | Filter by consistency_groups.luns.space.snapshot.space_used_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.used | integer | query | False | Filter by consistency_groups.luns.space.snapshot.used<br><br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.space.snapshot.reserve_size | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_size<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.autodelete.enabled | boolean | query | False | Filter by consistency_groups.luns.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.size | integer | query | False | Filter by consistency_groups.luns.space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |
| consistency_groups.luns.serial_number | string | query | False | Filter by consistency_groups.luns.serial_number<br>• maxLength: 12<br>• minLength: 12 |
| consistency_groups.luns.qos.policy.min_throughput_iops | integer | query | False | Filter by consistency_groups.luns.qos.policy.min_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.uuid | string | query | False | Filter by consistency_groups.luns.qos.policy.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.qos.policy.max_throughput_mbps | integer | query | False | Filter by consistency_groups.luns.qos.policy.max_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.name | string | query | False | Filter by consistency_groups.luns.qos.policy.name |
| consistency_groups.luns.qos.policy.min_throughput_mbps | integer | query | False | Filter by consistency_groups.luns.qos.policy.min_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.max_throughput_iops | integer | query | False | Filter by consistency_groups.luns.qos.policy.max_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| consistency_groups.luns.name | string | query | False | Filter by consistency_groups.luns.name |
| consistency_groups.luns.enabled | boolean | query | False | Filter by consistency_groups.luns.enabled |
| consistency_groups.luns.lun_maps.logical_unit_number | integer | query | False | Filter by consistency_groups.luns.lun_maps.logical_unit_number |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.lun_maps.igroup.igroups.uuid | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.igroups.uuid |
| consistency_groups.luns.lun_maps.igroup.igroups.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| consistency_groups.luns.lun_maps.igroup.uuid | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.uuid |
| consistency_groups.luns.lun_maps.igroup.protocol | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.protocol |
| consistency_groups.luns.lun_maps.igroup.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.name<br>• maxLength: 96<br>• minLength: 1 |
| consistency_groups.luns.lun_maps.igroup.os_type | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.os_type |
| consistency_groups.luns.lun_maps.igroup.comment | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.comment<br>• Introduced in: 9.11<br>• maxLength: 254<br>• minLength: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.lun_maps.igroup.initiators.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.initiators.name |
| consistency_groups.luns.lun_maps.igroup.initiators.comment | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.initiators.comment<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.tiering.policy | string | query | False | Filter by consistency_groups.tiering.policy |
| consistency_groups.application.type | string | query | False | Filter by consistency_groups.application.type<br>• Introduced in: 9.12 |
| consistency_groups.application.component_type | string | query | False | Filter by consistency_groups.application.component_type<br>• Introduced in: 9.12 |
| consistency_groups.name | string | query | False | Filter by consistency_groups.name |
| consistency_groups.parent_consistency_group.uuid | string | query | False | Filter by consistency_groups.parent_consistency_group.uuid |
| consistency_groups.parent_consistency_group.name | string | query | False | Filter by consistency_groups.parent_consistency_group.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.junction_parent.name | string | query | False | Filter by consistency_groups.volumes.nas.junction_parent.name<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.junction_parent.uuid | string | query | False | Filter by consistency_groups.volumes.nas.junction_parent.uuid<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.security_style | string | query | False | Filter by consistency_groups.volumes.nas.security_style<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.uid | integer | query | False | Filter by consistency_groups.volumes.nas.uid<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.gid | integer | query | False | Filter by consistency_groups.volumes.nas.gid<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.name | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.name<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| consistency_groups.volumes.nas.export_policy.rules.index | integer | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.index<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.clients.match | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.clients.match<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.chown_mode | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.chown_mode<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.ntfs_unix_security | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.ntfs_unix_security<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.allow_suid | boolean | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.allow_suid<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.export_policy.rules.superuser | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.superuser<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.ro_rule | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.ro_rule<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.protocols | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.protocols<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.anonymous_user | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.anonymous_user<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.allow_device_creation | boolean | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.allow_device_creation<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.export_policy.rules.rw_rule | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.rw_rule<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.id | integer | query | False | Filter by consistency_groups.volumes.nas.export_policy.id<br><br>• Introduced in: 9.14 |
| consistency_groups.volumes.nas.cifs.shares.namespace_caching | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.namespace_caching<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.continuously_available | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.continuously_available<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.vscan_profile | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.vscan_profile<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.type | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.type<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.acls.user_or_group | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.user_or_group<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.permission | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.permission<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.win_sid_unix_id | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.win_sid_unix_id<br><br>• Introduced in: 9.16 |
| consistency_groups.volumes.nas.cifs.shares.offline_files | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.offline_files<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.home_directory | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.home_directory<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.allow_unencrypted_access | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.allow_unencrypted_access<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.encryption | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.encryption<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.comment | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.comment<br>• Introduced in: 9.12<br>• maxLength: 256<br>• minLength: 1 |
| consistency_groups.volumes.nas.cifs.shares.dir_umask | integer | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.dir_umask<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.oplocks | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.oplocks<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.show_snapshot | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.show_snapshot<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.change_notify | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.change_notify<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.access_based_enumeration | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.access_based_enumeration<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.name | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.name<br>• Introduced in: 9.12<br>• maxLength: 80<br>• minLength: 1 |
| consistency_groups.volumes.nas.cifs.shares.file_umask | integer | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.file_umask<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.no_strict_security | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.no_strict_security<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.unix_symlink | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.unix_symlink<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.path | string | query | False | Filter by consistency_groups.volumes.nas.path<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.unix_permissions | integer | query | False | Filter by consistency_groups.volumes.nas.unix_permissions<br>• Introduced in: 9.12 |
| consistency_groups.volumes.comment | string | query | False | Filter by consistency_groups.volumes.comment<br>• maxLength: 1023<br>• minLength: 0 |
| consistency_groups.volumes.uuid | string | query | False | Filter by consistency_groups.volumes.uuid |
| consistency_groups.volumes.space.available | integer | query | False | Filter by consistency_groups.volumes.space.available |
| consistency_groups.volumes.space.size | integer | query | False | Filter by consistency_groups.volumes.space.size |
| consistency_groups.volumes.space.used | integer | query | False | Filter by consistency_groups.volumes.space.used |
| consistency_groups.volumes.snapshot_policy.uuid | string | query | False | Filter by consistency_groups.volumes.snapshot_policy.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.snapshot_policy.name | string | query | False | Filter by consistency_groups.volumes.snapshot_policy.name |
| consistency_groups.volumes.tiering.policy | string | query | False | Filter by consistency_groups.volumes.tiering.policy |
| consistency_groups.volumes.qos.policy.name | string | query | False | Filter by consistency_groups.volumes.qos.policy.name |
| consistency_groups.volumes.qos.policy.uuid | string | query | False | Filter by consistency_groups.volumes.qos.policy.uuid |
| consistency_groups.volumes.name | string | query | False | Filter by consistency_groups.volumes.name<br>• maxLength: 203<br>• minLength: 1 |
| consistency_groups._tags | string | query | False | Filter by consistency_groups._tags<br>• Introduced in: 9.15 |
| consistency_groups.qos.policy.name | string | query | False | Filter by consistency_groups.qos.policy.name |
| consistency_groups.qos.policy.uuid | string | query | False | Filter by consistency_groups.qos.policy.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.create_time | string | query | False | Filter by consistency_groups.namespaces.create_time<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.nsid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.nsid<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.tls.key_type | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.tls.key_type<br><br>• Introduced in: 9.16 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.nqn | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.nqn<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size<br><br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function<br><br>• Introduced in: 9.14 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode<br><br>• Introduced in: 9.16 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.priority | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.priority<br><br>• Introduced in: 9.14 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.local_svm | boolean | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.local_svm<br><br>• Introduced in: 9.17 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name<br><br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid<br>• Introduced in: 9.17 |
| consistency_groups.namespaces.subsystem_map.subsystem.uuid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.uuid<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.subsystem.comment | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.comment<br>• Introduced in: 9.12<br>• maxLength: 255<br>• minLength: 0 |
| consistency_groups.namespaces.subsystem_map.subsystem.os_type | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.os_type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.subsystem_map.subsystem.name | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.name<br>• Introduced in: 9.12<br>• maxLength: 64<br>• minLength: 1 |
| consistency_groups.namespaces.subsystem_map.anagrpid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.anagrpid<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.os_type | string | query | False | Filter by consistency_groups.namespaces.os_type<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.comment | string | query | False | Filter by consistency_groups.namespaces.comment<br>• Introduced in: 9.12<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.namespaces.uuid | string | query | False | Filter by consistency_groups.namespaces.uuid<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.snapshot.reserve_available | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_available<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.reserve_percent | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.space_used_percent | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.space_used_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.used | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.used<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.reserve_size | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_size<br><br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.snapshot.autodelete.enabled | boolean | query | False | Filter by consistency_groups.namespaces.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.size | integer | query | False | Filter by consistency_groups.namespaces.space.size<br>• Introduced in: 9.12<br>• Max value: 140737488355328<br>• Min value: 4096 |
| consistency_groups.namespaces.space.guarantee.reserved | boolean | query | False | Filter by consistency_groups.namespaces.space.guarantee.reserved<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.space.guarantee.requested | boolean | query | False | Filter by consistency_groups.namespaces.space.guarantee.requested<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.space.used | integer | query | False | Filter by consistency_groups.namespaces.space.used<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.block_size | integer | query | False | Filter by consistency_groups.namespaces.space.block_size<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.name | string | query | False | Filter by consistency_groups.namespaces.name<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.auto_delete | boolean | query | False | Filter by consistency_groups.namespaces.auto_delete<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.container_state | string | query | False | Filter by consistency_groups.namespaces.status.container_state<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.mapped | boolean | query | False | Filter by consistency_groups.namespaces.status.mapped<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.read_only | boolean | query | False | Filter by consistency_groups.namespaces.status.read_only<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.status.state | string | query | False | Filter by consistency_groups.namespaces.status.state<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.enabled | boolean | query | False | Filter by consistency_groups.namespaces.enabled<br><br>• Introduced in: 9.12 |
| consistency_groups.svm.name | string | query | False | Filter by consistency_groups.svm.name |
| consistency_groups.svm.uuid | string | query | False | Filter by consistency_groups.svm.uuid |
| uuid | string | query | False | Filter by uuid |
| replication_relationships.is_protected_by_svm_dr | boolean | query | False | Filter by replication_relationships.is_protected_by_svm_dr<br><br>• Introduced in: 9.14 |
| replication_relationships.uuid | string | query | False | Filter by replication_relationships.uuid<br><br>• Introduced in: 9.13 |
| replication_relationships.is_source | boolean | query | False | Filter by replication_relationships.is_source<br><br>• Introduced in: 9.13 |
| space.used | integer | query | False | Filter by space.used |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.size | integer | query | False | Filter by space.size |
| space.available | integer | query | False | Filter by space.available |
| qos.policy.name | string | query | False | Filter by qos.policy.name |
| qos.policy.uuid | string | query | False | Filter by qos.policy.uuid |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| namespaces.create_time | string | query | False | Filter by namespaces.create_time<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.nsid | string | query | False | Filter by namespaces.subsystem_map.nsid<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.subsystem.hosts.tls.key_type | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.tls.key_type<br>• Introduced in: 9.16 |
| namespaces.subsystem_map.subsystem.hosts.nqn | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.nqn<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode<br>• Introduced in: 9.16 |
| namespaces.subsystem_map.subsystem.hosts.priority | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.priority<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.proximity.local_svm | boolean | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.local_svm<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name<br>• Introduced in: 9.17 |
| namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid<br>• Introduced in: 9.17 |
| namespaces.subsystem_map.subsystem.uuid | string | query | False | Filter by namespaces.subsystem_map.subsystem.uuid<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.subsystem.comment | string | query | False | Filter by namespaces.subsystem_map.subsystem.comment<br>• Introduced in: 9.12<br>• maxLength: 255<br>• minLength: 0 |
| namespaces.subsystem_map.subsystem.os_type | string | query | False | Filter by namespaces.subsystem_map.subsystem.os_type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| namespaces.subsystem_map.subsystem.name | string | query | False | Filter by namespaces.subsystem_map.subsystem.name<br>• Introduced in: 9.12<br>• maxLength: 64<br>• minLength: 1 |
| namespaces.subsystem_map.anagrpid | string | query | False | Filter by namespaces.subsystem_map.anagrpid<br>• Introduced in: 9.12 |
| namespaces.os_type | string | query | False | Filter by namespaces.os_type<br>• Introduced in: 9.12 |
| namespaces.comment | string | query | False | Filter by namespaces.comment<br>• Introduced in: 9.12<br>• maxLength: 254<br>• minLength: 0 |
| namespaces.uuid | string | query | False | Filter by namespaces.uuid<br>• Introduced in: 9.12 |
| namespaces.space.snapshot.reserve_available | integer | query | False | Filter by namespaces.space.snapshot.reserve_available<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.space.snapshot.reserve_percent | integer | query | False | Filter by namespaces.space.snapshot.reserve_percent<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.space_used_percent | integer | query | False | Filter by namespaces.space.snapshot.space_used_percent<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.used | integer | query | False | Filter by namespaces.space.snapshot.used<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.reserve_size | integer | query | False | Filter by namespaces.space.snapshot.reserve_size<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.autodelete.enabled | boolean | query | False | Filter by namespaces.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| namespaces.space.size | integer | query | False | Filter by namespaces.space.size<br>• Introduced in: 9.12<br>• Max value: 140737488355328<br>• Min value: 4096 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.space.guarantee.reserved | boolean | query | False | Filter by namespaces.space.guarantee.reserved<br><br>• Introduced in: 9.12 |
| namespaces.space.guarantee.requested | boolean | query | False | Filter by namespaces.space.guarantee.requested<br><br>• Introduced in: 9.12 |
| namespaces.space.used | integer | query | False | Filter by namespaces.space.used<br><br>• Introduced in: 9.12 |
| namespaces.space.block_size | integer | query | False | Filter by namespaces.space.block_size<br><br>• Introduced in: 9.12 |
| namespaces.name | string | query | False | Filter by namespaces.name<br><br>• Introduced in: 9.12 |
| namespaces.auto_delete | boolean | query | False | Filter by namespaces.auto_delete<br><br>• Introduced in: 9.12 |
| namespaces.status.container_state | string | query | False | Filter by namespaces.status.container_state<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.status.mapped | boolean | query | False | Filter by namespaces.status.mapped<br><br>• Introduced in: 9.12 |
| namespaces.status.read_only | boolean | query | False | Filter by namespaces.status.read_only<br><br>• Introduced in: 9.12 |
| namespaces.status.state | string | query | False | Filter by namespaces.status.state<br><br>• Introduced in: 9.12 |
| namespaces.enabled | boolean | query | False | Filter by namespaces.enabled<br><br>• Introduced in: 9.12 |
| _tags | string | query | False | Filter by _tags<br><br>• Introduced in: 9.15 |
| volumes.nas.junction_parent.name | string | query | False | Filter by volumes.nas.junction_parent.name<br><br>• Introduced in: 9.12 |
| volumes.nas.junction_parent.uuid | string | query | False | Filter by volumes.nas.junction_parent.uuid<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.security_style | string | query | False | Filter by volumes.nas.security_style<br>• Introduced in: 9.12 |
| volumes.nas.uid | integer | query | False | Filter by volumes.nas.uid<br>• Introduced in: 9.12 |
| volumes.nas.gid | integer | query | False | Filter by volumes.nas.gid<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.name | string | query | False | Filter by volumes.nas.export_policy.name<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.index | integer | query | False | Filter by volumes.nas.export_policy.rules.index<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.clients.match | string | query | False | Filter by volumes.nas.export_policy.rules.clients.match<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.chown_mode | string | query | False | Filter by volumes.nas.export_policy.rules.chown_mode<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.export_policy.rules.ntfs_unix_security | string | query | False | Filter by volumes.nas.export_policy.rules.ntfs_unix_security<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.allow_suid | boolean | query | False | Filter by volumes.nas.export_policy.rules.allow_suid<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.superuser | string | query | False | Filter by volumes.nas.export_policy.rules.superuser<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.ro_rule | string | query | False | Filter by volumes.nas.export_policy.rules.ro_rule<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.protocols | string | query | False | Filter by volumes.nas.export_policy.rules.protocols<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.anonymous_user | string | query | False | Filter by volumes.nas.export_policy.rules.anonymous_user<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.export_policy.rules.allow_device_creation | boolean | query | False | Filter by volumes.nas.export_policy.rules.allow_device_creation<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.rw_rule | string | query | False | Filter by volumes.nas.export_policy.rules.rw_rule<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.id | integer | query | False | Filter by volumes.nas.export_policy.id<br>• Introduced in: 9.14 |
| volumes.nas.cifs.shares.namespace_caching | boolean | query | False | Filter by volumes.nas.cifs.shares.namespace_caching<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.continuously_available | boolean | query | False | Filter by volumes.nas.cifs.shares.continuously_available<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.vscan_profile | string | query | False | Filter by volumes.nas.cifs.shares.vscan_profile<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.acls.type | string | query | False | Filter by volumes.nas.cifs.shares.acls.type<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.user_or_group | string | query | False | Filter by volumes.nas.cifs.shares.acls.user_or_group<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.permission | string | query | False | Filter by volumes.nas.cifs.shares.acls.permission<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.win_sid_unix_id | string | query | False | Filter by volumes.nas.cifs.shares.acls.win_sid_unix_id<br><br>• Introduced in: 9.16 |
| volumes.nas.cifs.shares.offline_files | string | query | False | Filter by volumes.nas.cifs.shares.offline_files<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.home_directory | boolean | query | False | Filter by volumes.nas.cifs.shares.home_directory<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.allow_unencrypted_access | boolean | query | False | Filter by volumes.nas.cifs.shares.allow_unencrypted_access<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.encryption | boolean | query | False | Filter by volumes.nas.cifs.shares.encryption<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.comment | string | query | False | Filter by volumes.nas.cifs.shares.comment<br>• Introduced in: 9.12<br>• maxLength: 256<br>• minLength: 1 |
| volumes.nas.cifs.shares.dir_umask | integer | query | False | Filter by volumes.nas.cifs.shares.dir_umask<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.oplocks | boolean | query | False | Filter by volumes.nas.cifs.shares.oplocks<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.show_snapshot | boolean | query | False | Filter by volumes.nas.cifs.shares.show_snapshot<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.change_notify | boolean | query | False | Filter by volumes.nas.cifs.shares.change_notify<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.access_based_enumeration | boolean | query | False | Filter by volumes.nas.cifs.shares.access_based_enumeration<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.name | string | query | False | Filter by volumes.nas.cifs.shares.name<br>• Introduced in: 9.12<br>• maxLength: 80<br>• minLength: 1 |
| volumes.nas.cifs.shares.file_umask | integer | query | False | Filter by volumes.nas.cifs.shares.file_umask<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.no_strict_security | boolean | query | False | Filter by volumes.nas.cifs.shares.no_strict_security<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.unix_symlink | string | query | False | Filter by volumes.nas.cifs.shares.unix_symlink<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.path | string | query | False | Filter by volumes.nas.path<br>• Introduced in: 9.12 |
| volumes.nas.unix_permissions | integer | query | False | Filter by volumes.nas.unix_permissions<br>• Introduced in: 9.12 |
| volumes.comment | string | query | False | Filter by volumes.comment<br>• maxLength: 1023<br>• minLength: 0 |
| volumes.uuid | string | query | False | Filter by volumes.uuid |
| volumes.space.available | integer | query | False | Filter by volumes.space.available |
| volumes.space.size | integer | query | False | Filter by volumes.space.size |
| volumes.space.used | integer | query | False | Filter by volumes.space.used |
| volumes.snapshot_policy.uuid | string | query | False | Filter by volumes.snapshot_policy.uuid |
| volumes.snapshot_policy.name | string | query | False | Filter by volumes.snapshot_policy.name |
| volumes.tiering.policy | string | query | False | Filter by volumes.tiering.policy |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.qos.policy.name | string | query | False | Filter by volumes.qos.policy.name |
| volumes.qos.policy.uuid | string | query | False | Filter by volumes.qos.policy.uuid |
| volumes.name | string | query | False | Filter by volumes.name<br>• maxLength: 203<br>• minLength: 1 |
| parent_consistency_group.uuid | string | query | False | Filter by parent_consistency_group.uuid |
| parent_consistency_group.name | string | query | False | Filter by parent_consistency_group.name |
| vdisk_type | string | query | False | Filter by vdisk_type<br>• Introduced in: 9.17 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records. |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "application": {
        "component_type": "string",
        "type": "string"
      },
      "clone": {
        "guarantee": {
          "type": "string"
        },
        "parent_consistency_group": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "string",
          "uuid": "string"
        },
        "parent_snapshot": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "this_snapshot",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "parent_svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "split_complete_percent": 0,
        "split_estimate": 0,
        "volume": {
          "prefix": "string",
          "suffix": "string"
        }
      },
      "consistency_groups": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "_tags": [
            "team:csi",
            "environment:test"
          ],
          "application": {
            "component_type": "string",
            "type": "string"
          },
          "luns": [
            {
              "comment": "string",
              "create_time": "2018-06-04 19:00:00 +0000",
              "lun_maps": [
                {
                  "igroup": {
                    "comment": "string",
                    "igroups": [
                      {
                        "_links": {
                          "self": {
```

```
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "igroup1",
                    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
                  }
                ],
                "initiators": [
                  {
                    "comment": "my comment",
                    "name": "iqn.1998-01.com.corp.iscsi:name1"
                  }
                ],
                "name": "igroup1",
                "os_type": "string",
                "protocol": "string",
                "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
              }
            }
          ],
          "name": "/vol/volume1/lun1",
          "os_type": "string",
          "qos": {
            "policy": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "max_throughput_iops": 10000,
              "max_throughput_mbps": 500,
              "min_throughput_iops": 2000,
              "min_throughput_mbps": 500,
              "name": "performance",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "serial_number": "string",
          "space": {
            "size": 1073741824,
            "used": 0
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "name": "string",
```

```
          "namespaces": [
            {
              "comment": "string",
              "create_time": "2018-06-04 19:00:00 +0000",
              "name": "/vol/volume1/qtree1/namespace1",
              "os_type": "string",
              "space": {
                "block_size": 512,
                "size": 1073741824,
                "used": 0
              },
              "status": {
                "container_state": "string",
                "state": "online"
              },
              "subsystem_map": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "anagrpid": "00103050h",
                "nsid": "00000001h",
                "subsystem": {
                  "comment": "string",
                  "hosts": [
                    {
                      "dh_hmac_chap": {
                        "group_size": "string",
                        "hash_function": "string",
                        "mode": "bidirectional"
                      },
                      "nqn": "nqn.1992-01.example.com:string",
                      "priority": "string",
                      "proximity": {
                        "peer_svms": [
                          {
                            "_links": {
                              "self": {
                                "href": "/api/resourcelink"
                              }
                            },
                            "name": "peer1",
                            "uuid": "4204cf77-4c82-9bdb-5644-
b5a841c097a9"
                          }
```

```
                    ]
                  },
                  "tls": {
                    "key_type": "configured"
                  }
                }
              ],
              "name": "subsystem1",
              "os_type": "string",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "parent_consistency_group": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "my_consistency_group",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "qos": {
        "policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "restore_to": {
        "snapshot": {
          "name": "string",
          "uuid": "string"
        }
      },
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
      },
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "space": {
      "available": 5737418,
      "size": 1073741824,
      "used": 5737418
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "tiering": {
      "object_stores": [
        {}
      ],
      "policy": "string"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "volumes": [
      {
        "comment": "string",
        "name": "vol_cs_dept",
        "nas": {
          "cifs": {
            "shares": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "acls": [
                  {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "permission": "string",
```

```
                "type": "string",
                "user_or_group": "ENGDOMAIN\\ad_user",
                "win_sid_unix_id": "string"
              }
            ],
            "comment": "HR Department Share",
            "dir_umask": 18,
            "file_umask": 18,
            "name": "HR_SHARE",
            "offline_files": "string",
            "unix_symlink": "string",
            "vscan_profile": "string"
          }
        ]
      },
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 0,
        "name": "string",
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "anonymous_user": "string",
            "chown_mode": "string",
            "clients": [
              {
                "match": "0.0.0.0/0"
              }
            ],
            "ntfs_unix_security": "string",
            "protocols": [
              "string"
            ],
            "ro_rule": [
              "string"
            ],
            "rw_rule": [
              "string"
```

```
                ],
                "superuser": [
                  "string"
                ]
              }
            ]
          },
          "junction_parent": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "vs1_root",
            "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
          },
          "path": "/user/my_volume",
          "security_style": "string",
          "unix_permissions": 493
        },
        "qos": {
          "policy": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "snapshot_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "space": {
          "available": 0,
          "used": 0
        },
        "tiering": {
          "object_stores": [
```

```
                {}
              ],
              "policy": "string"
            },
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        ]
      }
    ],
    "luns": [
      {
        "comment": "string",
        "create_time": "2018-06-04 19:00:00 +0000",
        "lun_maps": [
          {
            "igroup": {
              "comment": "string",
              "igroups": [
                {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "name": "igroup1",
                  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
                }
              ],
              "initiators": [
                {
                  "comment": "my comment",
                  "name": "iqn.1998-01.com.corp.iscsi:name1"
                }
              ],
              "name": "igroup1",
              "os_type": "string",
              "protocol": "string",
              "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ],
        "name": "/vol/volume1/lun1",
        "os_type": "string",
        "qos": {
          "policy": {
            "_links": {
```

```
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "max_throughput_iops": 10000,
              "max_throughput_mbps": 500,
              "min_throughput_iops": 2000,
              "min_throughput_mbps": 500,
              "name": "performance",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "serial_number": "string",
          "space": {
            "size": 1073741824,
            "used": 0
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "available_space": 4096,
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "size": 4096,
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000",
```

```
        "used_space": 4096
      },
      "name": "string",
      "namespaces": [
        {
          "comment": "string",
          "create_time": "2018-06-04 19:00:00 +0000",
          "name": "/vol/volume1/qtree1/namespace1",
          "os_type": "string",
          "space": {
            "block_size": 512,
            "size": 1073741824,
            "used": 0
          },
          "status": {
            "container_state": "string",
            "state": "online"
          },
          "subsystem_map": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "anagrpid": "00103050h",
            "nsid": "00000001h",
            "subsystem": {
              "comment": "string",
              "hosts": [
                {
                  "dh_hmac_chap": {
                    "group_size": "string",
                    "hash_function": "string",
                    "mode": "bidirectional"
                  },
                  "nqn": "nqn.1992-01.example.com:string",
                  "priority": "string",
                  "proximity": {
                    "peer_svms": [
                      {
                        "_links": {
                          "self": {
                            "href": "/api/resourcelink"
                          }
                        },
                        "name": "peer1",
```

```
                        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                      }
                    ]
                  },
                  "tls": {
                    "key_type": "configured"
                  }
                }
              ],
              "name": "subsystem1",
              "os_type": "string",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "parent_consistency_group": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "my_consistency_group",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "qos": {
        "policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "replication_relationships": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
```

```
    ],
    "restore_to": {
      "snapshot": {
        "name": "string",
        "uuid": "string"
      }
    },
    "snapshot_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "space": {
      "available": 5737418,
      "size": 1073741824,
      "used": 5737418
    },
    "statistics": {
      "available_space": 4096,
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "size": 4096,
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "used_space": 4096
    },
    "svm": {
      "_links": {
        "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tiering": {
        "object_stores": [
          {}
        ],
        "policy": "string"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vdisk_type": "luns",
      "volumes": [
        {
          "comment": "string",
          "name": "vol_cs_dept",
          "nas": {
            "cifs": {
              "shares": [
                {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "acls": [
                    {
                      "_links": {
                        "self": {
                          "href": "/api/resourcelink"
                        }
                      },
                      "permission": "string",
                      "type": "string",
                      "user_or_group": "ENGDOMAIN\\ad_user",
                      "win_sid_unix_id": "string"
                    }
                  ],
                  "comment": "HR Department Share",
                  "dir_umask": 18,
                  "file_umask": 18,
                  "name": "HR_SHARE",
                  "offline_files": "string",
                  "unix_symlink": "string",
```

```
                "vscan_profile": "string"
              }
            ]
          },
          "export_policy": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "id": 0,
            "name": "string",
            "rules": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "anonymous_user": "string",
                "chown_mode": "string",
                "clients": [
                  {
                    "match": "0.0.0.0/0"
                  }
                ],
                "ntfs_unix_security": "string",
                "protocols": [
                  "string"
                ],
                "ro_rule": [
                  "string"
                ],
                "rw_rule": [
                  "string"
                ],
                "superuser": [
                  "string"
                ]
              }
            ]
          },
          "junction_parent": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
```

```
                }
              },
              "name": "vs1_root",
              "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
            },
            "path": "/user/my_volume",
            "security_style": "string",
            "unix_permissions": 493
          },
          "qos": {
            "policy": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "performance",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "snapshot_policy": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "default",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "space": {
            "available": 0,
            "used": 0
          },
          "tiering": {
            "object_stores": [
              {}
            ],
            "policy": "string"
          },
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ]
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| component_type | string | Nested consistency group tag. |
| type | string | Top level consistency group tag. |

guarantee

| Name | Type | Description |
|---|---|---|
| type | string | The type of space guarantee of this volume in the aggregate. |

parent_consistency_group

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

parent_snapshot

Consistency group that is to be cloned.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume name suffix/prefix for the cloned volumes.

| Name | Type | Description |
| --- | --- | --- |
| prefix | string | Volume name prefix for cloned volumes. |
| suffix | string | Volume name suffix for cloned volumes. |

clone

Creates a clone of an existing consistency group from the current contents or an existing snapshot.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | |

| Name | Type | Description |
| --- | --- | --- |
| is_flexclone | boolean | Specifies if this consistency group is a FlexClone of a consistency group. |
| parent_consistency_group | parent_consistency_group | Consistency group that is to be cloned. |
| parent_snapshot | parent_snapshot | Consistency group that is to be cloned. |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| split_complete_percent | integer | Percentage of FlexClone blocks split from its parent consistency group. |
| split_estimate | integer | Space required to split the FlexClone consistency group. |
| split_initiated | boolean | Splits volumes after cloning. Defaults to false during POST. Only accepts true during a PATCH. |
| volume | volume | Volume name suffix/prefix for the cloned volumes. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

igroups

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br><br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not reduced using the REST interface.<br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform, and the available space in the containing volume and aggregate.<br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

luns

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the LUN was created. |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the `state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| lun_maps | array[lun_maps] | An array of LUN maps.<br><br>A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups. |
| name | string | The fully qualified path name of the LUN composed of the "/vol" prefix, the volume name, the qtree name (optional), and the base name of the LUN. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| qos | qos | |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br><br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br><br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace, in bytes.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not reduced using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |

| Name | Type | Description |
| --- | --- | --- |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

consistency_group_nvme_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

consistency_group_nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[consistency_group_nvme_host] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`.
They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | consistency_group_nvme_subsystem | An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts. |

namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |
| space | space | The storage space related properties of the NVMe namespace. |
| status | status | Status information about the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`.<br>They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. |
| uuid | string | The unique identifier of the NVMe namespace. |

parent_consistency_group

The parent consistency group.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

storage_service

Determines the placement of any storage object created during this operation.

| Name | Type | Description |
|---|---|---|
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| name | string | New name for consistency group. Required to resolve naming collisions. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot

A consistency group's snapshot

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group's snapshot to restore to. |
| uuid | string | The UUID of the consistency group's snapshot to restore to. |

restore_to

Use to restore a consistency group to a previous snapshot

| Name | Type | Description |
| --- | --- | --- |
| snapshot | snapshot | A consistency group's snapshot |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

space

Space information for the consistency group.

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in the consistency group, in bytes. |
| size | integer | The total provisioned size of the consistency group, in bytes. |
| used | integer | The amount of space consumed in the consistency group, in bytes. |

svm

The Storage Virtual Machine (SVM) in which the consistency group is located.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

object_stores

tiering

The tiering placement and policy definitions for volumes in this consistency group.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |
| win_sid_unix_id | string | Windows SID/UNIX ID depending on access-control type. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value:<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value:<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| shares | array[consistency_group_cifs_share] | |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |

| Name | Type | Description |
| --- | --- | --- |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

export_policy

The policy associated with volumes to export them for protocol access.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| id | integer | Identifier for the export policy. |
| name | string | Name of the export policy. |
| rules | array[export_rules] | The set of rules that govern the export policy. |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

The CIFS share policy and/or export policies for this volume.

| Name | Type | Description |
| --- | --- | --- |
| cifs | cifs | |
| export_policy | export_policy | The policy associated with volumes to export them for protocol access. |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within an SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ UNIX-style security. |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number, consisting of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). First digit selects the set user ID (4), set group ID (2), and sticky (1) attributes. Second digit selects permission for the owner of the file. Third selects permissions for other users in the same group while the fourth selects permissions for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group, and other permissions are given (as in 755, representing the second, third and fourth digit), the first digit is assumed to be zero. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

qos

The QoS policy for this volume.

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available space, in bytes. |
| size | integer | Total provisioned size, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |

tiering

The tiering placement and policy definitions for this volume.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment for the volume. Valid in POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | The CIFS share policy and/or export policies for this volume. |
| qos | qos | The QoS policy for this volume. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | |
| tiering | tiering | The tiering placement and policy definitions for this volume. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |

consistency_groups

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |

| Name | Type | Description |
| --- | --- | --- |
| application | application | |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| qos | qos | |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | Space information for the consistency group. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency observed at the storage object, in microseconds.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| available_space | integer | The total space available in the consistency group, in bytes. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency observed at the storage object, in microseconds. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance and capacity data. |
| used_space | integer | The total space used in the consistency group, in bytes. |

replication_relationships

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| is_protected_by_svm_dr | boolean | Indicates whether or not this consistency group is protected by SVM DR. |
| is_source | boolean | Indicates whether or not this consistency group is the source for replication. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SnapMirror relationship. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| available_space | integer | The total space available in the consistency group, in bytes. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |
| used_space | integer | The total used space in the consistency group, in bytes. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

records

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |
| metric | metric | Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | Space information for the consistency group. |
| statistics | statistics | These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| vdisk_type | string | Type of objects in the consistency group. |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a consistency group

POST `/application/consistency-groups`

**Introduced In:** 9.10

Creates a consistency group with one or more consistency groups having:

- new SAN volumes,
- existing SAN, NVMe or NAS FlexVol volumes in a new or existing consistency group

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the group.
- `volumes`
- `luns` or `namespaces`

#### Naming Conventions

#### Consistency groups

- name or consistency_groups[].name, if specified
- derived from volumes[0].name, if only one volume is specified, same as volume name

#### Volume

- volumes[].name, if specified
- derived from volume prefix in luns[].name
- derived from cg[].name, suffixed by "_#" where "#" is a system generated unique number
- suffixed by "_#" where "#" is a system generated unique number, if provisioning_options.count is provided

#### LUN

- luns[].name, if specified
- derived from volumes[].name, suffixed by "_#" where "#" is a system generated unique number
- suffixed by "_#" where "#" is a system generated unique number, if provisioning_options.count is provided

#### NVMe Namespace

- namespaces[].name, if specified
- derived from volumes[].name, suffixed by "_#" where "#" is a system generated unique number
- suffixed by "_#" where "#" is a system generated unique number, if provisioning_options.count is provided

#### Related ONTAP commands

- `vserver consistency-group create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group. The maximum number of items for this array is 16. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| vdisk_type | string | Type of objects in the consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

**Example request**

```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "application": {
    "component_type": "string",
    "type": "string"
  },
  "clone": {
    "guarantee": {
      "type": "string"
    },
    "parent_consistency_group": {
      "name": "string",
      "uuid": "string"
    },
    "parent_snapshot": {
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0,
    "volume": {
      "prefix": "string",
      "suffix": "string"
    }
  },
  "consistency_groups": [
    {
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "application": {
        "component_type": "string",
        "type": "string"
      },
      "luns": [
        {
```

```
        "clone": {
          "source": {
            "name": "/vol/volume1/lun1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "comment": "string",
        "create_time": "2018-06-04 19:00:00 +0000",
        "lun_maps": [
          {
            "igroup": {
              "comment": "string",
              "igroups": [
                {
                  "name": "igroup1",
                  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
                }
              ],
              "initiators": [
                {
                  "comment": "my comment",
                  "name": "iqn.1998-01.com.corp.iscsi:name1"
                }
              ],
              "name": "igroup1",
              "os_type": "string",
              "protocol": "string",
              "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ],
        "name": "/vol/volume1/lun1",
        "os_type": "string",
        "provisioning_options": {
          "action": "string"
        },
        "qos": {
          "policy": {
            "max_throughput_iops": 10000,
            "max_throughput_mbps": 500,
            "min_throughput_iops": 2000,
            "min_throughput_mbps": 500,
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
```

```
          "serial_number": "string",
          "space": {
            "size": 1073741824,
            "used": 0
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "name": "string",
      "namespaces": [
        {
          "comment": "string",
          "create_time": "2018-06-04 19:00:00 +0000",
          "name": "/vol/volume1/qtree1/namespace1",
          "os_type": "string",
          "provisioning_options": {
            "action": "string"
          },
          "space": {
            "block_size": 512,
            "size": 1073741824,
            "used": 0
          },
          "subsystem_map": {
            "anagrpid": "00103050h",
            "nsid": "00000001h",
            "subsystem": {
              "comment": "string",
              "hosts": [
                {
                  "dh_hmac_chap": {
                    "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
                    "group_size": "string",
                    "hash_function": "string",
                    "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
                    "mode": "bidirectional"
                  },
                  "nqn": "nqn.1992-01.example.com:string",
                  "priority": "string",
                  "proximity": {
                    "peer_svms": [
                      {
                        "name": "peer1",
                        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
```

```
                      }
                    ]
                  },
                  "tls": {
                    "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
                    "key_type": "configured"
                  }
                }
              ],
              "name": "subsystem1",
              "os_type": "string",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "parent_consistency_group": {
        "name": "my_consistency_group",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "provisioning_options": {
        "action": "string",
        "name": "string",
        "storage_service": {
          "name": "string"
        }
      },
      "qos": {
        "policy": {
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "restore_to": {
        "snapshot": {
          "name": "string",
          "uuid": "string"
        }
      },
      "snapshot_policy": {
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
```

```
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tiering": {
        "control": "string",
        "object_stores": [
          {
            "name": "string"
          }
        ],
        "policy": "string"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volumes": [
        {
          "comment": "string",
          "name": "vol_cs_dept",
          "nas": {
            "cifs": {
              "shares": [
                {
                  "acls": [
                    {
                      "permission": "string",
                      "type": "string",
                      "user_or_group": "ENGDOMAIN\\ad_user"
                    }
                  ],
                  "comment": "HR Department Share",
                  "dir_umask": 18,
                  "file_umask": 18,
                  "name": "HR_SHARE",
                  "offline_files": "string",
                  "unix_symlink": "string",
                  "vscan_profile": "string"
                }
              ]
            },
            "export_policy": {
              "id": 0,
              "name": "string",
              "rules": [
                {
                  "anonymous_user": "string",
                  "chown_mode": "string",
                  "clients": [
```

```
                    {
                      "match": "0.0.0.0/0"
                    }
                  ],
                  "ntfs_unix_security": "string",
                  "protocols": [
                    "string"
                  ],
                  "ro_rule": [
                    "string"
                  ],
                  "rw_rule": [
                    "string"
                  ],
                  "superuser": [
                    "string"
                  ]
                }
              ]
            },
            "junction_parent": {
              "name": "vs1_root",
              "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
            },
            "path": "/user/my_volume",
            "security_style": "string",
            "unix_permissions": 493
          },
          "provisioning_options": {
            "action": "string",
            "storage_service": {
              "name": "string"
            }
          },
          "qos": {
            "policy": {
              "name": "performance",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "snapshot_policy": {
            "name": "default",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "space": {
            "available": 0,
```

```
            "used": 0
          },
          "tiering": {
            "control": "string",
            "object_stores": [
              {
                "name": "string"
              }
            ],
            "policy": "string"
          },
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ]
    }
  ],
  "luns": [
    {
      "clone": {
        "source": {
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "lun_maps": [
        {
          "igroup": {
            "comment": "string",
            "igroups": [
              {
                "name": "igroup1",
                "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
              }
            ],
            "initiators": [
              {
                "comment": "my comment",
                "name": "iqn.1998-01.com.corp.iscsi:name1"
              }
            ],
            "name": "igroup1",
            "os_type": "string",
            "protocol": "string",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
```

```
          }
        }
      ],
      "name": "/vol/volume1/lun1",
      "os_type": "string",
      "provisioning_options": {
        "action": "string"
      },
      "qos": {
        "policy": {
          "max_throughput_iops": 10000,
          "max_throughput_mbps": 500,
          "min_throughput_iops": 2000,
          "min_throughput_mbps": 500,
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "serial_number": "string",
      "space": {
        "size": 1073741824,
        "used": 0
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "string",
  "namespaces": [
    {
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "name": "/vol/volume1/qtree1/namespace1",
      "os_type": "string",
      "provisioning_options": {
        "action": "string"
      },
      "space": {
        "block_size": 512,
        "size": 1073741824,
        "used": 0
      },
      "subsystem_map": {
        "anagrpid": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "comment": "string",
```

```
          "hosts": [
            {
              "dh_hmac_chap": {
                "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
                "group_size": "string",
                "hash_function": "string",
                "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
                "mode": "bidirectional"
              },
              "nqn": "nqn.1992-01.example.com:string",
              "priority": "string",
              "proximity": {
                "peer_svms": [
                  {
                    "name": "peer1",
                    "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                  }
                ]
              },
              "tls": {
                "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
                "key_type": "configured"
              }
            }
          ],
          "name": "subsystem1",
          "os_type": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "parent_consistency_group": {
    "name": "my_consistency_group",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "provisioning_options": {
    "action": "string",
    "name": "string",
    "storage_service": {
      "name": "string"
    }
```

```
  },
  "qos": {
    "policy": {
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "replication_relationships": [
    {
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "restore_to": {
    "snapshot": {
      "name": "string",
      "uuid": "string"
    }
  },
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "control": "string",
    "object_stores": [
      {
        "name": "string"
      }
    ],
    "policy": "string"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vdisk_type": "luns",
  "volumes": [
    {
      "comment": "string",
      "name": "vol_cs_dept",
      "nas": {
        "cifs": {
          "shares": [
            {
              "acls": [
```

```
                {
                  "permission": "string",
                  "type": "string",
                  "user_or_group": "ENGDOMAIN\\ad_user"
                }
              ],
              "comment": "HR Department Share",
              "dir_umask": 18,
              "file_umask": 18,
              "name": "HR_SHARE",
              "offline_files": "string",
              "unix_symlink": "string",
              "vscan_profile": "string"
            }
          ]
        },
        "export_policy": {
          "id": 0,
          "name": "string",
          "rules": [
            {
              "anonymous_user": "string",
              "chown_mode": "string",
              "clients": [
                {
                  "match": "0.0.0.0/0"
                }
              ],
              "ntfs_unix_security": "string",
              "protocols": [
                "string"
              ],
              "ro_rule": [
                "string"
              ],
              "rw_rule": [
                "string"
              ],
              "superuser": [
                "string"
              ]
            }
          ]
        },
        "junction_parent": {
          "name": "vs1_root",
```

```
            "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
          },
          "path": "/user/my_volume",
          "security_style": "string",
          "unix_permissions": 493
        },
        "provisioning_options": {
          "action": "string",
          "storage_service": {
            "name": "string"
          }
        },
        "qos": {
          "policy": {
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "snapshot_policy": {
          "name": "default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "space": {
          "available": 0,
          "used": 0
        },
        "tiering": {
          "control": "string",
          "object_stores": [
            {
              "name": "string"
            }
          ],
          "policy": "string"
        },
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ]
  }
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374127 | The specified LUN name is invalid. |
| 5376461 | The specified name is invalid because it does not match the format requirements. |
| 5440501 | The specified storage unit size is insufficient. |
| 5440509 | No suitable storage can be found for the specified requirements. |
| 53411842 | Consistency group does not exist. |
| 53411843 | A consistency group with specified UUID was not found. |

| Error Code | Description |
| --- | --- |
| 53411844 | Specified consistency group was not found in the specified SVM. |
| 53411845 | The specified UUID and name refer to different consistency groups. |
| 53411846 | Either name or UUID must be provided. |
| 53411853 | Fields provided in the request conflict with each other. |
| 53411856 | The provided field is only supported when provisioning new objects. |
| 53411860 | An object with the same identifier in the same scope exists. |
| 53411861 | The specified volume does not exist in provided volume array. |
| 53411862 | Modifying existing igroups is not supported using this API. |
| 53411864 | Request content insufficient to add an existing volume to an application. |
| 53411865 | Volumes contained in one consistency group cannot be added to a different consistency group. |
| 53411866 | LUNs are not supported on FlexGroup volumes. |
| 53411867 | LUN name is too long after appending a unique suffix. |
| 53411869 | Volume name is too long after appending a unique suffix. |
| 53411870 | When using the "round_robin" layout, the volume count must not be greater than the LUN count. |
| 53411959 | Volumes with snapshot locking enabled cannot be added to a consistency group. |
| 53412040 | Splitting a non-SnapLock clone from a Snaplock consistency group during clone creation is not supported. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

application

| Name | Type | Description |
|---|---|---|
| component_type | string | Nested consistency group tag. |
| type | string | Top level consistency group tag. |

guarantee

| Name | Type | Description |
|---|---|---|
| type | string | The type of space guarantee of this volume in the aggregate. |

parent_consistency_group

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

parent_snapshot

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume name suffix/prefix for the cloned volumes.

| Name | Type | Description |
| --- | --- | --- |
| prefix | string | Volume name prefix for cloned volumes. |
| suffix | string | Volume name suffix for cloned volumes. |

clone

Creates a clone of an existing consistency group from the current contents or an existing snapshot.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | |
| is_flexclone | boolean | Specifies if this consistency group is a FlexClone of a consistency group. |
| parent_consistency_group | parent_consistency_group | Consistency group that is to be cloned. |
| parent_snapshot | parent_snapshot | Consistency group that is to be cloned. |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| split_complete_percent | integer | Percentage of FlexClone blocks split from its parent consistency group. |
| split_estimate | integer | Space required to split the FlexClone consistency group. |

| Name | Type | Description |
|---|---|---|
| split_initiated | boolean | Splits volumes after cloning. Defaults to false during POST. Only accepts true during a PATCH. |
| volume | volume | Volume name suffix/prefix for the cloned volumes. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |
| uuid | string | The unique identifier of the clone source LUN.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

| Name | Type | Description |
| --- | --- | --- |
| source | source | The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.<br><br>Valid in POST to create a new LUN as a clone of the source.<br><br>Valid in PATCH to overwrite an existing LUN's data as a clone of another. |

igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br><br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not reduced using the REST interface.<br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform, and the available space in the containing volume and aggregate.<br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

luns

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the `state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| lun_maps | array[lun_maps] | An array of LUN maps.<br><br>A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups. |
| name | string | The fully qualified path name of the LUN composed of the "/vol" prefix, the volume name, the qtree name (optional), and the base name of the LUN. Valid in POST and PATCH. |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace, in bytes.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not reduced using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |

| Name | Type | Description |
| --- | --- | --- |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| controller_secret_key | string | The controller secret for NVMe in-band authentication. The value of this property is used by the NVMe host to authenticate the NVMe controller while establishing a connection. If unset, the controller is not authenticated. When supplied, the property `host_secret_key` must also be supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a controller secret has been set for the host, but the controller secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| host_secret_key | string | The host secret for NVMe in-band authentication. The value of this property is used by the NVMe controller to authenticate the NVMe host while establishing a connection. If unset, no authentication is performed by the host or controller. This property must be supplied if any other NVMe in-band authentication properties are supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a host secret has been set for the host, but the host secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |

| Name | Type | Description |
| --- | --- | --- |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| configured_psk | string | A user supplied pre-shared key (PSK) value in PSK Interchange Format. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree. This property is only allowed when `key_type` is `configured`. If `configured_psk` is supplied and `key_type` is unset, `key_type` defaults to `configured`.<br><br>This property is write-only. The `key_type` property can be used to identify if a configured PSK has been set for the host, but the PSK value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

consistency_group_nvme_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

consistency_group_nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| hosts | array[consistency_group_nvme_host] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | consistency_group_nvme_subsystem | An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts. |

namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
|---|---|---|
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| space | space | The storage space related properties of the NVMe namespace. |

| Name | Type | Description |
| --- | --- | --- |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`.<br>They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. |
| uuid | string | The unique identifier of the NVMe namespace. |

parent_consistency_group

The parent consistency group.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

storage_service

Determines the placement of any storage object created during this operation.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |

| Name | Type | Description |
| --- | --- | --- |
| name | string | New name for consistency group. Required to resolve naming collisions. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

policy

The QoS policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot

A consistency group's snapshot

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group's snapshot to restore to. |
| uuid | string | The UUID of the consistency group's snapshot to restore to. |

restore_to

Use to restore a consistency group to a previous snapshot

| Name | Type | Description |
| --- | --- | --- |
| snapshot | snapshot | A consistency group's snapshot |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

space

Space information for the consistency group.

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in the consistency group, in bytes. |
| size | integer | The total provisioned size of the consistency group, in bytes. |
| used | integer | The amount of space consumed in the consistency group, in bytes. |

svm

The Storage Virtual Machine (SVM) in which the consistency group is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for volumes in this consistency group.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |

| Name | Type | Description |
| --- | --- | --- |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |

| Name | Type | Description |
| --- | --- | --- |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |

| Name | Type | Description |
| --- | --- | --- |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| shares | array[consistency_group_cifs_share] | |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

export_policy

The policy associated with volumes to export them for protocol access.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Identifier for the export policy. |
| name | string | Name of the export policy. |
| rules | array[export_rules] | The set of rules that govern the export policy. |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the parent volume. |

nas

The CIFS share policy and/or export policies for this volume.

| Name | Type | Description |
| --- | --- | --- |
| cifs | cifs | |
| export_policy | export_policy | The policy associated with volumes to export them for protocol access. |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within an SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ UNIX-style security. |

| Name | Type | Description |
| --- | --- | --- |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number, consisting of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). First digit selects the set user ID (4), set group ID (2), and sticky (1) attributes. Second digit selects permission for the owner of the file. Third selects permissions for other users in the same group while the fourth selects permissions for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group, and other permissions are given (as in 755, representing the second, third and fourth digit), the first digit is assumed to be zero. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

qos

The QoS policy for this volume.

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available space, in bytes. |
| size | integer | Total provisioned size, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |

tiering

The tiering placement and policy definitions for this volume.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment for the volume. Valid in POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | The CIFS share policy and/or export policies for this volume. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | The QoS policy for this volume. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | |
| tiering | tiering | The tiering placement and policy definitions for this volume. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |

consistency_groups

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency observed at the storage object, in microseconds.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space.

| Name | Type | Description |
| --- | --- | --- |
| available_space | integer | The total space available in the consistency group, in bytes. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency observed at the storage object, in microseconds. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance and capacity data. |
| used_space | integer | The total space used in the consistency group, in bytes. |

replication_relationships

| Name | Type | Description |
| --- | --- | --- |
| is_protected_by_svm_dr | boolean | Indicates whether or not this consistency group is protected by SVM DR. |
| is_source | boolean | Indicates whether or not this consistency group is the source for replication. |
| uuid | string | The unique identifier of the SnapMirror relationship. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| available_space | integer | The total space available in the consistency group, in bytes. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |
| used_space | integer | The total used space in the consistency group, in bytes. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

consistency_group

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |

| Name | Type | Description |
| --- | --- | --- |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| vdisk_type | string | Type of objects in the consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |

| Name | Type | Description |
| --- | --- | --- |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |

| Name | Type | Description |
| --- | --- | --- |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a consistency group

DELETE `/application/consistency-groups/{uuid}`

**Introduced In:** 9.10

Deletes a consistency group.

> this will not delete any associated volumes or LUNs. To delete those elements, use the appropriate object endpoint.

#### Related ONTAP commands

- `vserver consistency-group delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the consistency group to delete. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53411842 | Consistency group does not exist. |
| 53411843 | A consistency group with specified UUID was not found. |
| 53411844 | Specified consistency group was not found in the specified SVM. |
| 53411845 | The specified UUID and name refer to different consistency groups. |
| 53411846 | Either name or UUID must be provided. |
| 53412041 | Cannot delete a consistency group because it contains one or more storage units that are part of replication relationships. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a consistency group

GET `/application/consistency-groups/{uuid}`

**Introduced In:** 9.10

Retrieves a single consistency group.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See DOC Requesting specific fields to learn more.

- `volumes`
- `luns`
- `namespaces`

#### Related ONTAP commands

- `vserver consistency-group show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the group to retrieve. |
| application.component_type | string | query | False | Filter by application.component_type<br>• Introduced in: 9.12 |
| application.type | string | query | False | Filter by application.type<br>• Introduced in: 9.12 |
| tiering.policy | string | query | False | Filter by tiering.policy |
| luns.create_time | string | query | False | Filter by luns.create_time |
| luns.comment | string | query | False | Filter by luns.comment<br>• maxLength: 254<br>• minLength: 0 |
| luns.os_type | string | query | False | Filter by luns.os_type |
| luns.uuid | string | query | False | Filter by luns.uuid |
| luns.space.used | integer | query | False | Filter by luns.space.used |
| luns.space.guarantee.requested | boolean | query | False | Filter by luns.space.guarantee.requested<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.space.guarantee.reserved | boolean | query | False | Filter by luns.space.guarantee.reserved<br><br>• Introduced in: 9.11 |
| luns.space.snapshot.reserve_available | integer | query | False | Filter by luns.space.snapshot.reserve_available<br><br>• Introduced in: 9.18 |
| luns.space.snapshot.reserve_percent | integer | query | False | Filter by luns.space.snapshot.reserve_percent<br><br>• Introduced in: 9.18 |
| luns.space.snapshot.space_used_percent | integer | query | False | Filter by luns.space.snapshot.space_used_percent<br><br>• Introduced in: 9.18 |
| luns.space.snapshot.used | integer | query | False | Filter by luns.space.snapshot.used<br><br>• Introduced in: 9.18 |
| luns.space.snapshot.reserve_size | integer | query | False | Filter by luns.space.snapshot.reserve_size<br><br>• Introduced in: 9.18 |
| luns.space.snapshot.autodelete.enabled | boolean | query | False | Filter by luns.space.snapshot.autodelete.enabled<br><br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.space.size | integer | query | False | Filter by luns.space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |
| luns.serial_number | string | query | False | Filter by luns.serial_number<br>• maxLength: 12<br>• minLength: 12 |
| luns.qos.policy.min_throughput_iops | integer | query | False | Filter by luns.qos.policy.min_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| luns.qos.policy.uuid | string | query | False | Filter by luns.qos.policy.uuid |
| luns.qos.policy.max_throughput_mbps | integer | query | False | Filter by luns.qos.policy.max_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| luns.qos.policy.name | string | query | False | Filter by luns.qos.policy.name |
| luns.qos.policy.min_throughput_mbps | integer | query | False | Filter by luns.qos.policy.min_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.qos.policy.max_throughput_iops | integer | query | False | Filter by luns.qos.policy.max_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| luns.name | string | query | False | Filter by luns.name |
| luns.enabled | boolean | query | False | Filter by luns.enabled |
| luns.lun_maps.logical_unit_number | integer | query | False | Filter by luns.lun_maps.logical_unit_number |
| luns.lun_maps.igroup.igroups.uuid | string | query | False | Filter by luns.lun_maps.igroup.igroups.uuid |
| luns.lun_maps.igroup.igroups.name | string | query | False | Filter by luns.lun_maps.igroup.igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| luns.lun_maps.igroup.uuid | string | query | False | Filter by luns.lun_maps.igroup.uuid |
| luns.lun_maps.igroup.protocol | string | query | False | Filter by luns.lun_maps.igroup.protocol |
| luns.lun_maps.igroup.name | string | query | False | Filter by luns.lun_maps.igroup.name<br>• maxLength: 96<br>• minLength: 1 |
| luns.lun_maps.igroup.os_type | string | query | False | Filter by luns.lun_maps.igroup.os_type |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| luns.lun_maps.igroup.comment | string | query | False | Filter by luns.lun_maps.igroup.comment<br>• Introduced in: 9.11<br>• maxLength: 254<br>• minLength: 0 |
| luns.lun_maps.igroup.initiators.name | string | query | False | Filter by luns.lun_maps.igroup.initiators.name |
| luns.lun_maps.igroup.initiators.comment | string | query | False | Filter by luns.lun_maps.igroup.initiators.comment<br>• maxLength: 254<br>• minLength: 0 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.13 |
| statistics.used_space | integer | query | False | Filter by statistics.used_space<br>• Introduced in: 9.13 |
| statistics.throughput_raw.other | integer | query | False | Filter by statistics.throughput_raw.other<br>• Introduced in: 9.13 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.13 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.13 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.13 |
| statistics.available_space | integer | query | False | Filter by statistics.available_space<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.13 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.13 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.13 |
| statistics.size | integer | query | False | Filter by statistics.size<br><br>• Introduced in: 9.13 |
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid |
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| replication_source | boolean | query | False | Filter by replication_source |
| name | string | query | False | Filter by name |
| replicated | boolean | query | False | Filter by replicated |
| clone.guarantee.type | string | query | False | Filter by clone.guarantee.type<br>• Introduced in: 9.12 |
| clone.parent_snapshot.name | string | query | False | Filter by clone.parent_snapshot.name<br>• Introduced in: 9.12 |
| clone.parent_snapshot.uuid | string | query | False | Filter by clone.parent_snapshot.uuid<br>• Introduced in: 9.15 |
| clone.parent_svm.name | string | query | False | Filter by clone.parent_svm.name<br>• Introduced in: 9.15 |
| clone.parent_svm.uuid | string | query | False | Filter by clone.parent_svm.uuid<br>• Introduced in: 9.15 |
| clone.volume.suffix | string | query | False | Filter by clone.volume.suffix<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| clone.volume.prefix | string | query | False | Filter by clone.volume.prefix<br><br>• Introduced in: 9.12 |
| clone.split_complete_percent | integer | query | False | Filter by clone.split_complete_percent<br><br>• Introduced in: 9.15 |
| clone.parent_consistency_group.name | string | query | False | Filter by clone.parent_consistency_group.name<br><br>• Introduced in: 9.12 |
| clone.parent_consistency_group.uuid | string | query | False | Filter by clone.parent_consistency_group.uuid<br><br>• Introduced in: 9.12 |
| clone.is_flexclone | boolean | query | False | Filter by clone.is_flexclone<br><br>• Introduced in: 9.15 |
| clone.split_initiated | boolean | query | False | Filter by clone.split_initiated<br><br>• Introduced in: 9.12 |
| clone.split_estimate | integer | query | False | Filter by clone.split_estimate<br><br>• Introduced in: 9.15 |
| metric.size | integer | query | False | Filter by metric.size<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.13 |
| metric.available_space | integer | query | False | Filter by metric.available_space<br><br>• Introduced in: 9.13 |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.13 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.13 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.13 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.13 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.13 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.13 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.13 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.13 |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other<br><br>• Introduced in: 9.13 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.13 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.13 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.timestamp | string | query | False | Filter by metric.timestamp<br>• Introduced in: 9.13 |
| metric.used_space | integer | query | False | Filter by metric.used_space<br>• Introduced in: 9.13 |
| consistency_groups.uuid | string | query | False | Filter by consistency_groups.uuid |
| consistency_groups.space.used | integer | query | False | Filter by consistency_groups.space.used |
| consistency_groups.space.size | integer | query | False | Filter by consistency_groups.space.size |
| consistency_groups.space.available | integer | query | False | Filter by consistency_groups.space.available |
| consistency_groups.snapshot_policy.uuid | string | query | False | Filter by consistency_groups.snapshot_policy.uuid |
| consistency_groups.snapshot_policy.name | string | query | False | Filter by consistency_groups.snapshot_policy.name |
| consistency_groups.luns.create_time | string | query | False | Filter by consistency_groups.luns.create_time |
| consistency_groups.luns.comment | string | query | False | Filter by consistency_groups.luns.comment<br>• maxLength: 254<br>• minLength: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.os_type | string | query | False | Filter by consistency_groups.luns.os_type |
| consistency_groups.luns.uuid | string | query | False | Filter by consistency_groups.luns.uuid |
| consistency_groups.luns.space.used | integer | query | False | Filter by consistency_groups.luns.space.used |
| consistency_groups.luns.space.guarantee.requested | boolean | query | False | Filter by consistency_groups.luns.space.guarantee.requested<br>• Introduced in: 9.11 |
| consistency_groups.luns.space.guarantee.reserved | boolean | query | False | Filter by consistency_groups.luns.space.guarantee.reserved<br>• Introduced in: 9.11 |
| consistency_groups.luns.space.snapshot.reserve_available | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_available<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.reserve_percent | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_percent<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.space.snapshot.space_used_percent | integer | query | False | Filter by consistency_groups.luns.space.snapshot.space_used_percent<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.used | integer | query | False | Filter by consistency_groups.luns.space.snapshot.used<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.reserve_size | integer | query | False | Filter by consistency_groups.luns.space.snapshot.reserve_size<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.snapshot.autodelete.enabled | boolean | query | False | Filter by consistency_groups.luns.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| consistency_groups.luns.space.size | integer | query | False | Filter by consistency_groups.luns.space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |
| consistency_groups.luns.serial_number | string | query | False | Filter by consistency_groups.luns.serial_number<br>• maxLength: 12<br>• minLength: 12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.qos.policy.min_throughput_iops | integer | query | False | Filter by consistency_groups.luns.qos.policy.min_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.uuid | string | query | False | Filter by consistency_groups.luns.qos.policy.uuid |
| consistency_groups.luns.qos.policy.max_throughput_mbps | integer | query | False | Filter by consistency_groups.luns.qos.policy.max_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.name | string | query | False | Filter by consistency_groups.luns.qos.policy.name |
| consistency_groups.luns.qos.policy.min_throughput_mbps | integer | query | False | Filter by consistency_groups.luns.qos.policy.min_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| consistency_groups.luns.qos.policy.max_throughput_iops | integer | query | False | Filter by consistency_groups.luns.qos.policy.max_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.name | string | query | False | Filter by consistency_groups.luns.name |
| consistency_groups.luns.enabled | boolean | query | False | Filter by consistency_groups.luns.enabled |
| consistency_groups.luns.lun_maps.logical_unit_number | integer | query | False | Filter by consistency_groups.luns.lun_maps.logical_unit_number |
| consistency_groups.luns.lun_maps.igroup.igroups.uuid | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.igroups.uuid |
| consistency_groups.luns.lun_maps.igroup.igroups.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| consistency_groups.luns.lun_maps.igroup.uuid | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.uuid |
| consistency_groups.luns.lun_maps.igroup.protocol | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.protocol |
| consistency_groups.luns.lun_maps.igroup.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.name<br>• maxLength: 96<br>• minLength: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.luns.lun_maps.igroup.os_type | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.os_type |
| consistency_groups.luns.lun_maps.igroup.comment | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.comment<br>• Introduced in: 9.11<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.luns.lun_maps.igroup.initiators.name | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.initiators.name |
| consistency_groups.luns.lun_maps.igroup.initiators.comment | string | query | False | Filter by consistency_groups.luns.lun_maps.igroup.initiators.comment<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.tiering.policy | string | query | False | Filter by consistency_groups.tiering.policy |
| consistency_groups.application.type | string | query | False | Filter by consistency_groups.application.type<br>• Introduced in: 9.12 |
| consistency_groups.application.component_type | string | query | False | Filter by consistency_groups.application.component_type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.name | string | query | False | Filter by consistency_groups.name |
| consistency_groups.parent_consistency_group.uuid | string | query | False | Filter by consistency_groups.parent_consistency_group.uuid |
| consistency_groups.parent_consistency_group.name | string | query | False | Filter by consistency_groups.parent_consistency_group.name |
| consistency_groups.volumes.nas.junction_parent.name | string | query | False | Filter by consistency_groups.volumes.nas.junction_parent.name<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.junction_parent.uuid | string | query | False | Filter by consistency_groups.volumes.nas.junction_parent.uuid<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.security_style | string | query | False | Filter by consistency_groups.volumes.nas.security_style<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.uid | integer | query | False | Filter by consistency_groups.volumes.nas.uid<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.gid | integer | query | False | Filter by consistency_groups.volumes.nas.gid<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.name | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.name<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.index | integer | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.index<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.clients.match | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.clients.match<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.chown_mode | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.chown_mode<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.ntfs_unix_security | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.ntfs_unix_security<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| consistency_groups.volumes.nas.export_policy.rules.allow_suid | boolean | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.allow_suid<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.superuser | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.superuser<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.ro_rule | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.ro_rule<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.protocols | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.protocols<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.anonymous_user | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.anonymous_user<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.export_policy.rules.allow_device_creation | boolean | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.allow_device_creation<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.rules.rw_rule | string | query | False | Filter by consistency_groups.volumes.nas.export_policy.rules.rw_rule<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.export_policy.id | integer | query | False | Filter by consistency_groups.volumes.nas.export_policy.id<br><br>• Introduced in: 9.14 |
| consistency_groups.volumes.nas.cifs.shares.namespace_caching | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.namespace_caching<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.continuously_available | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.continuously_available<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.vscan_profile | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.vscan_profile<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.acls.type | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.type<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.user_or_group | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.user_or_group<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.permission | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.permission<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.acls.win_sid_unix_id | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.acls.win_sid_unix_id<br><br>• Introduced in: 9.16 |
| consistency_groups.volumes.nas.cifs.shares.offline_files | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.offline_files<br><br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.home_directory | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.home_directory<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.allow_unencrypted_access | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.allow_unencrypted_access<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.encryption | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.encryption<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.comment | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.comment<br>• Introduced in: 9.12<br>• maxLength: 256<br>• minLength: 1 |
| consistency_groups.volumes.nas.cifs.shares.dir_umask | integer | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.dir_umask<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.oplocks | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.oplocks<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.show_snapshot | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.show_snapshot<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.change_notify | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.change_notify<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.access_based_enumeration | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.access_based_enumeration<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.name | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.name<br>• Introduced in: 9.12<br>• maxLength: 80<br>• minLength: 1 |
| consistency_groups.volumes.nas.cifs.shares.file_umask | integer | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.file_umask<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.nas.cifs.shares.no_strict_security | boolean | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.no_strict_security<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.cifs.shares.unix_symlink | string | query | False | Filter by consistency_groups.volumes.nas.cifs.shares.unix_symlink<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.path | string | query | False | Filter by consistency_groups.volumes.nas.path<br>• Introduced in: 9.12 |
| consistency_groups.volumes.nas.unix_permissions | integer | query | False | Filter by consistency_groups.volumes.nas.unix_permissions<br>• Introduced in: 9.12 |
| consistency_groups.volumes.comment | string | query | False | Filter by consistency_groups.volumes.comment<br>• maxLength: 1023<br>• minLength: 0 |
| consistency_groups.volumes.uuid | string | query | False | Filter by consistency_groups.volumes.uuid |
| consistency_groups.volumes.space.available | integer | query | False | Filter by consistency_groups.volumes.space.available |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.volumes.space.size | integer | query | False | Filter by consistency_groups.volumes.space.size |
| consistency_groups.volumes.space.used | integer | query | False | Filter by consistency_groups.volumes.space.used |
| consistency_groups.volumes.snapshot_policy.uuid | string | query | False | Filter by consistency_groups.volumes.snapshot_policy.uuid |
| consistency_groups.volumes.snapshot_policy.name | string | query | False | Filter by consistency_groups.volumes.snapshot_policy.name |
| consistency_groups.volumes.tiering.policy | string | query | False | Filter by consistency_groups.volumes.tiering.policy |
| consistency_groups.volumes.qos.policy.name | string | query | False | Filter by consistency_groups.volumes.qos.policy.name |
| consistency_groups.volumes.qos.policy.uuid | string | query | False | Filter by consistency_groups.volumes.qos.policy.uuid |
| consistency_groups.volumes.name | string | query | False | Filter by consistency_groups.volumes.name<br>• maxLength: 203<br>• minLength: 1 |
| consistency_groups._tags | string | query | False | Filter by consistency_groups._tags<br>• Introduced in: 9.15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.qos.policy.name | string | query | False | Filter by consistency_groups.qos.policy.name |
| consistency_groups.qos.policy.uuid | string | query | False | Filter by consistency_groups.qos.policy.uuid |
| consistency_groups.namespaces.create_time | string | query | False | Filter by consistency_groups.namespaces.create_time<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.nsid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.nsid<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.tls.key_type | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.tls.key_type<br>• Introduced in: 9.16 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.nqn | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.nqn<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size<br><br>• Introduced in: 9.14 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function<br><br>• Introduced in: 9.14 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode<br><br>• Introduced in: 9.16 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.priority | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.priority<br><br>• Introduced in: 9.14 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.local_svm | boolean | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.local_svm<br><br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name<br>• Introduced in: 9.17 |
| consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid<br>• Introduced in: 9.17 |
| consistency_groups.namespaces.subsystem_map.subsystem.uuid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.uuid<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.subsystem_map.subsystem.comment | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.comment<br>• Introduced in: 9.12<br>• maxLength: 255<br>• minLength: 0 |
| consistency_groups.namespaces.subsystem_map.subsystem.os_type | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.os_type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| consistency_groups.namespaces.subsystem_map.subsystem.name | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.subsystem.name<br>• Introduced in: 9.12<br>• maxLength: 64<br>• minLength: 1 |
| consistency_groups.namespaces.subsystem_map.anagrpid | string | query | False | Filter by consistency_groups.namespaces.subsystem_map.anagrpid<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.os_type | string | query | False | Filter by consistency_groups.namespaces.os_type<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.comment | string | query | False | Filter by consistency_groups.namespaces.comment<br>• Introduced in: 9.12<br>• maxLength: 254<br>• minLength: 0 |
| consistency_groups.namespaces.uuid | string | query | False | Filter by consistency_groups.namespaces.uuid<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.snapshot.reserve_available | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_available<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.reserve_percent | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.space_used_percent | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.space_used_percent<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.used | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.used<br><br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.snapshot.reserve_size | integer | query | False | Filter by consistency_groups.namespaces.space.snapshot.reserve_size<br><br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.snapshot.autodelete.enabled | boolean | query | False | Filter by consistency_groups.namespaces.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| consistency_groups.namespaces.space.size | integer | query | False | Filter by consistency_groups.namespaces.space.size<br>• Introduced in: 9.12<br>• Max value: 140737488355328<br>• Min value: 4096 |
| consistency_groups.namespaces.space.guarantee.reserved | boolean | query | False | Filter by consistency_groups.namespaces.space.guarantee.reserved<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.space.guarantee.requested | boolean | query | False | Filter by consistency_groups.namespaces.space.guarantee.requested<br>• Introduced in: 9.12 |
| consistency_groups.namespaces.space.used | integer | query | False | Filter by consistency_groups.namespaces.space.used<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.space.block_size | integer | query | False | Filter by consistency_groups.namespaces.space.block_size<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.name | string | query | False | Filter by consistency_groups.namespaces.name<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.auto_delete | boolean | query | False | Filter by consistency_groups.namespaces.auto_delete<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.container_state | string | query | False | Filter by consistency_groups.namespaces.status.container_state<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.mapped | boolean | query | False | Filter by consistency_groups.namespaces.status.mapped<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.status.read_only | boolean | query | False | Filter by consistency_groups.namespaces.status.read_only<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_groups.namespaces.status.state | string | query | False | Filter by consistency_groups.namespaces.status.state<br><br>• Introduced in: 9.12 |
| consistency_groups.namespaces.enabled | boolean | query | False | Filter by consistency_groups.namespaces.enabled<br><br>• Introduced in: 9.12 |
| consistency_groups.svm.name | string | query | False | Filter by consistency_groups.svm.name |
| consistency_groups.svm.uuid | string | query | False | Filter by consistency_groups.svm.uuid |
| replication_relationships.is_protected_by_svm_dr | boolean | query | False | Filter by replication_relationships.is_protected_by_svm_dr<br><br>• Introduced in: 9.14 |
| replication_relationships.uuid | string | query | False | Filter by replication_relationships.uuid<br><br>• Introduced in: 9.13 |
| replication_relationships.is_source | boolean | query | False | Filter by replication_relationships.is_source<br><br>• Introduced in: 9.13 |
| space.used | integer | query | False | Filter by space.used |
| space.size | integer | query | False | Filter by space.size |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.available | integer | query | False | Filter by space.available |
| qos.policy.name | string | query | False | Filter by qos.policy.name |
| qos.policy.uuid | string | query | False | Filter by qos.policy.uuid |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| namespaces.create_time | string | query | False | Filter by namespaces.create_time<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.nsid | string | query | False | Filter by namespaces.subsystem_map.nsid<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.subsystem.hosts.tls.key_type | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.tls.key_type<br>• Introduced in: 9.16 |
| namespaces.subsystem_map.subsystem.hosts.nqn | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.nqn<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.group_size<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.dh_hmac_chap.mode<br>• Introduced in: 9.16 |
| namespaces.subsystem_map.subsystem.hosts.priority | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.priority<br>• Introduced in: 9.14 |
| namespaces.subsystem_map.subsystem.hosts.proximity.local_svm | boolean | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.local_svm<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.name<br>• Introduced in: 9.17 |
| namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid | string | query | False | Filter by namespaces.subsystem_map.subsystem.hosts.proximity.peer_svms.uuid<br>• Introduced in: 9.17 |
| namespaces.subsystem_map.subsystem.uuid | string | query | False | Filter by namespaces.subsystem_map.subsystem.uuid<br>• Introduced in: 9.12 |
| namespaces.subsystem_map.subsystem.comment | string | query | False | Filter by namespaces.subsystem_map.subsystem.comment<br>• Introduced in: 9.12<br>• maxLength: 255<br>• minLength: 0 |
| namespaces.subsystem_map.subsystem.os_type | string | query | False | Filter by namespaces.subsystem_map.subsystem.os_type<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.subsystem_map.subsystem.name | string | query | False | Filter by namespaces.subsystem_map.subsystem.name<br>• Introduced in: 9.12<br>• maxLength: 64<br>• minLength: 1 |
| namespaces.subsystem_map.anagrpid | string | query | False | Filter by namespaces.subsystem_map.anagrpid<br>• Introduced in: 9.12 |
| namespaces.os_type | string | query | False | Filter by namespaces.os_type<br>• Introduced in: 9.12 |
| namespaces.comment | string | query | False | Filter by namespaces.comment<br>• Introduced in: 9.12<br>• maxLength: 254<br>• minLength: 0 |
| namespaces.uuid | string | query | False | Filter by namespaces.uuid<br>• Introduced in: 9.12 |
| namespaces.space.snapshot.reserve_available | integer | query | False | Filter by namespaces.space.snapshot.reserve_available<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| namespaces.space.snapshot.reserve_percent | integer | query | False | Filter by namespaces.space.snapshot.reserve_percent<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.space_used_percent | integer | query | False | Filter by namespaces.space.snapshot.space_used_percent<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.used | integer | query | False | Filter by namespaces.space.snapshot.used<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.reserve_size | integer | query | False | Filter by namespaces.space.snapshot.reserve_size<br>• Introduced in: 9.18 |
| namespaces.space.snapshot.autodelete.enabled | boolean | query | False | Filter by namespaces.space.snapshot.autodelete.enabled<br>• Introduced in: 9.18 |
| namespaces.space.size | integer | query | False | Filter by namespaces.space.size<br>• Introduced in: 9.12<br>• Max value: 140737488355328<br>• Min value: 4096 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.space.guarantee.reserved | boolean | query | False | Filter by namespaces.space.guarantee.reserved<br><br>• Introduced in: 9.12 |
| namespaces.space.guarantee.requested | boolean | query | False | Filter by namespaces.space.guarantee.requested<br><br>• Introduced in: 9.12 |
| namespaces.space.used | integer | query | False | Filter by namespaces.space.used<br><br>• Introduced in: 9.12 |
| namespaces.space.block_size | integer | query | False | Filter by namespaces.space.block_size<br><br>• Introduced in: 9.12 |
| namespaces.name | string | query | False | Filter by namespaces.name<br><br>• Introduced in: 9.12 |
| namespaces.auto_delete | boolean | query | False | Filter by namespaces.auto_delete<br><br>• Introduced in: 9.12 |
| namespaces.status.container_state | string | query | False | Filter by namespaces.status.container_state<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| namespaces.status.mapped | boolean | query | False | Filter by namespaces.status.mapped<br><br>• Introduced in: 9.12 |
| namespaces.status.read_only | boolean | query | False | Filter by namespaces.status.read_only<br><br>• Introduced in: 9.12 |
| namespaces.status.state | string | query | False | Filter by namespaces.status.state<br><br>• Introduced in: 9.12 |
| namespaces.enabled | boolean | query | False | Filter by namespaces.enabled<br><br>• Introduced in: 9.12 |
| _tags | string | query | False | Filter by _tags<br><br>• Introduced in: 9.15 |
| volumes.nas.junction_parent.name | string | query | False | Filter by volumes.nas.junction_parent.name<br><br>• Introduced in: 9.12 |
| volumes.nas.junction_parent.uuid | string | query | False | Filter by volumes.nas.junction_parent.uuid<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.security_style | string | query | False | Filter by volumes.nas.security_style<br>• Introduced in: 9.12 |
| volumes.nas.uid | integer | query | False | Filter by volumes.nas.uid<br>• Introduced in: 9.12 |
| volumes.nas.gid | integer | query | False | Filter by volumes.nas.gid<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.name | string | query | False | Filter by volumes.nas.export_policy.name<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.index | integer | query | False | Filter by volumes.nas.export_policy.rules.index<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.clients.match | string | query | False | Filter by volumes.nas.export_policy.rules.clients.match<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.chown_mode | string | query | False | Filter by volumes.nas.export_policy.rules.chown_mode<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.export_policy.rules.ntfs_unix_security | string | query | False | Filter by volumes.nas.export_policy.rules.ntfs_unix_security<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.allow_suid | boolean | query | False | Filter by volumes.nas.export_policy.rules.allow_suid<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.superuser | string | query | False | Filter by volumes.nas.export_policy.rules.superuser<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.ro_rule | string | query | False | Filter by volumes.nas.export_policy.rules.ro_rule<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.protocols | string | query | False | Filter by volumes.nas.export_policy.rules.protocols<br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.anonymous_user | string | query | False | Filter by volumes.nas.export_policy.rules.anonymous_user<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.export_policy.rules.allow_device_creation | boolean | query | False | Filter by volumes.nas.export_policy.rules.allow_device_creation<br><br>• Introduced in: 9.12 |
| volumes.nas.export_policy.rules.rw_rule | string | query | False | Filter by volumes.nas.export_policy.rules.rw_rule<br><br>• Introduced in: 9.12 |
| volumes.nas.export_policy.id | integer | query | False | Filter by volumes.nas.export_policy.id<br><br>• Introduced in: 9.14 |
| volumes.nas.cifs.shares.namespace_caching | boolean | query | False | Filter by volumes.nas.cifs.shares.namespace_caching<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.continuously_available | boolean | query | False | Filter by volumes.nas.cifs.shares.continuously_available<br><br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.vscan_profile | string | query | False | Filter by volumes.nas.cifs.shares.vscan_profile<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.acls.type | string | query | False | Filter by volumes.nas.cifs.shares.acls.type<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.user_or_group | string | query | False | Filter by volumes.nas.cifs.shares.acls.user_or_group<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.permission | string | query | False | Filter by volumes.nas.cifs.shares.acls.permission<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.acls.win_sid_unix_id | string | query | False | Filter by volumes.nas.cifs.shares.acls.win_sid_unix_id<br>• Introduced in: 9.16 |
| volumes.nas.cifs.shares.offline_files | string | query | False | Filter by volumes.nas.cifs.shares.offline_files<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.home_directory | boolean | query | False | Filter by volumes.nas.cifs.shares.home_directory<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.allow_unencrypted_access | boolean | query | False | Filter by volumes.nas.cifs.shares.allow_unencrypted_access<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.encryption | boolean | query | False | Filter by volumes.nas.cifs.shares.encryption<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.comment | string | query | False | Filter by volumes.nas.cifs.shares.comment<br>• Introduced in: 9.12<br>• maxLength: 256<br>• minLength: 1 |
| volumes.nas.cifs.shares.dir_umask | integer | query | False | Filter by volumes.nas.cifs.shares.dir_umask<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.oplocks | boolean | query | False | Filter by volumes.nas.cifs.shares.oplocks<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.show_snapshot | boolean | query | False | Filter by volumes.nas.cifs.shares.show_snapshot<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.cifs.shares.change_notify | boolean | query | False | Filter by volumes.nas.cifs.shares.change_notify<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.access_based_enumeration | boolean | query | False | Filter by volumes.nas.cifs.shares.access_based_enumeration<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.name | string | query | False | Filter by volumes.nas.cifs.shares.name<br>• Introduced in: 9.12<br>• maxLength: 80<br>• minLength: 1 |
| volumes.nas.cifs.shares.file_umask | integer | query | False | Filter by volumes.nas.cifs.shares.file_umask<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.no_strict_security | boolean | query | False | Filter by volumes.nas.cifs.shares.no_strict_security<br>• Introduced in: 9.12 |
| volumes.nas.cifs.shares.unix_symlink | string | query | False | Filter by volumes.nas.cifs.shares.unix_symlink<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.nas.path | string | query | False | Filter by volumes.nas.path<br>• Introduced in: 9.12 |
| volumes.nas.unix_permissions | integer | query | False | Filter by volumes.nas.unix_permissions<br>• Introduced in: 9.12 |
| volumes.comment | string | query | False | Filter by volumes.comment<br>• maxLength: 1023<br>• minLength: 0 |
| volumes.uuid | string | query | False | Filter by volumes.uuid |
| volumes.space.available | integer | query | False | Filter by volumes.space.available |
| volumes.space.size | integer | query | False | Filter by volumes.space.size |
| volumes.space.used | integer | query | False | Filter by volumes.space.used |
| volumes.snapshot_policy.uuid | string | query | False | Filter by volumes.snapshot_policy.uuid |
| volumes.snapshot_policy.name | string | query | False | Filter by volumes.snapshot_policy.name |
| volumes.tiering.policy | string | query | False | Filter by volumes.tiering.policy |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volumes.qos.policy.name | string | query | False | Filter by volumes.qos.policy.name |
| volumes.qos.policy.uuid | string | query | False | Filter by volumes.qos.policy.uuid |
| volumes.name | string | query | False | Filter by volumes.name<br>• maxLength: 203<br>• minLength: 1 |
| parent_consistency_group.uuid | string | query | False | Filter by parent_consistency_group.uuid |
| parent_consistency_group.name | string | query | False | Filter by parent_consistency_group.name |
| vdisk_type | string | query | False | Filter by vdisk_type<br>• Introduced in: 9.17 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group. The maximum number of items for this array is 16. |
| metric | metric | Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| qos | qos | |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | Space information for the consistency group. |
| statistics | statistics | These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vdisk_type | string | Type of objects in the consistency group. |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "application": {
    "component_type": "string",
    "type": "string"
  },
  "clone": {
    "guarantee": {
      "type": "string"
    },
    "parent_consistency_group": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "string"
    },
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
      },
      "split_complete_percent": 0,
      "split_estimate": 0,
      "volume": {
        "prefix": "string",
        "suffix": "string"
      }
    },
    "consistency_groups": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "_tags": [
          "team:csi",
          "environment:test"
        ],
        "application": {
          "component_type": "string",
          "type": "string"
        },
        "luns": [
          {
            "comment": "string",
            "create_time": "2018-06-04 19:00:00 +0000",
            "lun_maps": [
              {
                "igroup": {
                  "comment": "string",
                  "igroups": [
                    {
                      "_links": {
                        "self": {
                          "href": "/api/resourcelink"
                        }
                      },
                      "name": "igroup1",
                      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
                    }
                  ],
                  "initiators": [
                    {
                      "comment": "my comment",
                      "name": "iqn.1998-01.com.corp.iscsi:name1"
```

```
                }
              ],
              "name": "igroup1",
              "os_type": "string",
              "protocol": "string",
              "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ],
        "name": "/vol/volume1/lun1",
        "os_type": "string",
        "qos": {
          "policy": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "max_throughput_iops": 10000,
            "max_throughput_mbps": 500,
            "min_throughput_iops": 2000,
            "min_throughput_mbps": 500,
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "serial_number": "string",
        "space": {
          "size": 1073741824,
          "used": 0
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "name": "string",
    "namespaces": [
      {
        "comment": "string",
        "create_time": "2018-06-04 19:00:00 +0000",
        "name": "/vol/volume1/qtree1/namespace1",
        "os_type": "string",
        "space": {
          "block_size": 512,
          "size": 1073741824,
          "used": 0
        },
```

```
      "status": {
        "container_state": "string",
        "state": "online"
      },
      "subsystem_map": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "anagrpid": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "comment": "string",
          "hosts": [
            {
              "dh_hmac_chap": {
                "group_size": "string",
                "hash_function": "string",
                "mode": "bidirectional"
              },
              "nqn": "nqn.1992-01.example.com:string",
              "priority": "string",
              "proximity": {
                "peer_svms": [
                  {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "peer1",
                    "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                  }
                ]
              },
              "tls": {
                "key_type": "configured"
              }
            }
          ],
          "name": "subsystem1",
          "os_type": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "parent_consistency_group": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "my_consistency_group",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "qos": {
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "restore_to": {
      "snapshot": {
        "name": "string",
        "uuid": "string"
      }
    },
    "snapshot_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "space": {
      "available": 5737418,
      "size": 1073741824,
      "used": 5737418
    },
    "svm": {
      "_links": {
        "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tiering": {
        "object_stores": [
          {}
        ],
        "policy": "string"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volumes": [
        {
          "comment": "string",
          "name": "vol_cs_dept",
          "nas": {
            "cifs": {
              "shares": [
                {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "acls": [
                    {
                      "_links": {
                        "self": {
                          "href": "/api/resourcelink"
                        }
                      },
                      "permission": "string",
                      "type": "string",
                      "user_or_group": "ENGDOMAIN\\ad_user",
                      "win_sid_unix_id": "string"
                    }
                  ],
                  "comment": "HR Department Share",
                  "dir_umask": 18,
                  "file_umask": 18,
                  "name": "HR_SHARE",
                  "offline_files": "string",
                  "unix_symlink": "string",
                  "vscan_profile": "string"
```

```
            }
          ]
        },
        "export_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "id": 0,
          "name": "string",
          "rules": [
            {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "anonymous_user": "string",
              "chown_mode": "string",
              "clients": [
                {
                  "match": "0.0.0.0/0"
                }
              ],
              "ntfs_unix_security": "string",
              "protocols": [
                "string"
              ],
              "ro_rule": [
                "string"
              ],
              "rw_rule": [
                "string"
              ],
              "superuser": [
                "string"
              ]
            }
          ]
        },
        "junction_parent": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
```

```
            },
            "name": "vs1_root",
            "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
          },
          "path": "/user/my_volume",
          "security_style": "string",
          "unix_permissions": 493
        },
        "qos": {
          "policy": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "snapshot_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "space": {
          "available": 0,
          "used": 0
        },
        "tiering": {
          "object_stores": [
            {}
          ],
          "policy": "string"
        },
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ]
  }
],
"luns": [
  {
    "comment": "string",
```

```
      "create_time": "2018-06-04 19:00:00 +0000",
      "lun_maps": [
        {
          "igroup": {
            "comment": "string",
            "igroups": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "igroup1",
                "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
              }
            ],
            "initiators": [
              {
                "comment": "my comment",
                "name": "iqn.1998-01.com.corp.iscsi:name1"
              }
            ],
            "name": "igroup1",
            "os_type": "string",
            "protocol": "string",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "name": "/vol/volume1/lun1",
      "os_type": "string",
      "qos": {
        "policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "max_throughput_iops": 10000,
          "max_throughput_mbps": 500,
          "min_throughput_iops": 2000,
          "min_throughput_mbps": 500,
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
```

```
      "serial_number": "string",
      "space": {
        "size": 1073741824,
        "used": 0
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "available_space": 4096,
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "size": 4096,
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000",
    "used_space": 4096
  },
  "name": "string",
  "namespaces": [
    {
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "name": "/vol/volume1/qtree1/namespace1",
      "os_type": "string",
      "space": {
        "block_size": 512,
        "size": 1073741824,
```

```
        "used": 0
      },
      "status": {
        "container_state": "string",
        "state": "online"
      },
      "subsystem_map": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "anagrpid": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "comment": "string",
          "hosts": [
            {
              "dh_hmac_chap": {
                "group_size": "string",
                "hash_function": "string",
                "mode": "bidirectional"
              },
              "nqn": "nqn.1992-01.example.com:string",
              "priority": "string",
              "proximity": {
                "peer_svms": [
                  {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "peer1",
                    "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                  }
                ]
              },
              "tls": {
                "key_type": "configured"
              }
            }
          ],
          "name": "subsystem1",
          "os_type": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "parent_consistency_group": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "my_consistency_group",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "replication_relationships": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "restore_to": {
    "snapshot": {
      "name": "string",
      "uuid": "string"
    }
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
```

```
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "available": 5737418,
    "size": 1073741824,
    "used": 5737418
  },
  "statistics": {
    "available_space": 4096,
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "size": 4096,
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000",
    "used_space": 4096
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_stores": [
      {}
    ],
    "policy": "string"
  },
```

```
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vdisk_type": "luns",
  "volumes": [
    {
      "comment": "string",
      "name": "vol_cs_dept",
      "nas": {
        "cifs": {
          "shares": [
            {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "acls": [
                {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "permission": "string",
                  "type": "string",
                  "user_or_group": "ENGDOMAIN\\ad_user",
                  "win_sid_unix_id": "string"
                }
              ],
              "comment": "HR Department Share",
              "dir_umask": 18,
              "file_umask": 18,
              "name": "HR_SHARE",
              "offline_files": "string",
              "unix_symlink": "string",
              "vscan_profile": "string"
            }
          ]
        },
        "export_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "id": 0,
          "name": "string",
```

```
          "rules": [
            {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "anonymous_user": "string",
              "chown_mode": "string",
              "clients": [
                {
                  "match": "0.0.0.0/0"
                }
              ],
              "ntfs_unix_security": "string",
              "protocols": [
                "string"
              ],
              "ro_rule": [
                "string"
              ],
              "rw_rule": [
                "string"
              ],
              "superuser": [
                "string"
              ]
            }
          ]
        },
        "junction_parent": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "vs1_root",
          "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
        },
        "path": "/user/my_volume",
        "security_style": "string",
        "unix_permissions": 493
      },
      "qos": {
        "policy": {
          "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "space": {
        "available": 0,
        "used": 0
      },
      "tiering": {
        "object_stores": [
          {}
        ],
        "policy": "string"
      },
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53411842 | Consistency group does not exist. |
| 53411843 | A consistency group with specified UUID was not found. |

| Error Code | Description |
| --- | --- |
| 53411844 | Specified consistency group was not found in the specified SVM. |
| 53411845 | The specified UUID and name refer to different consistency groups. |
| 53411846 | Either name or UUID must be provided. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

application

| Name | Type | Description |
|---|---|---|
| component_type | string | Nested consistency group tag. |
| type | string | Top level consistency group tag. |

guarantee

| Name | Type | Description |
|---|---|---|
| type | string | The type of space guarantee of this volume in the aggregate. |

parent_consistency_group

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

parent_snapshot

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume name suffix/prefix for the cloned volumes.

| Name | Type | Description |
|---|---|---|
| prefix | string | Volume name prefix for cloned volumes. |
| suffix | string | Volume name suffix for cloned volumes. |

clone

Creates a clone of an existing consistency group from the current contents or an existing snapshot.

| Name | Type | Description |
|---|---|---|
| guarantee | guarantee | |
| is_flexclone | boolean | Specifies if this consistency group is a FlexClone of a consistency group. |
| parent_consistency_group | parent_consistency_group | Consistency group that is to be cloned. |

| Name | Type | Description |
| --- | --- | --- |
| parent_snapshot | parent_snapshot | Consistency group that is to be cloned. |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| split_complete_percent | integer | Percentage of FlexClone blocks split from its parent consistency group. |
| split_estimate | integer | Space required to split the FlexClone consistency group. |
| split_initiated | boolean | Splits volumes after cloning. Defaults to false during POST. Only accepts true during a PATCH. |
| volume | volume | Volume name suffix/prefix for the cloned volumes. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

igroups

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br><br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not reduced using the REST interface.<br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform, and the available space in the containing volume and aggregate.<br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

luns

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the LUN was created. |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the `state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| lun_maps | array[lun_maps] | An array of LUN maps.<br><br>A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups. |
| name | string | The fully qualified path name of the LUN composed of the "/vol" prefix, the volume name, the qtree name (optional), and the base name of the LUN. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| qos | qos | |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br><br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br><br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace, in bytes.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not reduced using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |

| Name | Type | Description |
| --- | --- | --- |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

consistency_group_nvme_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

consistency_group_nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[consistency_group_nvme_host] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`.
They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | consistency_group_nvme_subsystem | An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts. |

namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |
| space | space | The storage space related properties of the NVMe namespace. |
| status | status | Status information about the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`.<br>They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. |
| uuid | string | The unique identifier of the NVMe namespace. |

parent_consistency_group

The parent consistency group.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

storage_service

Determines the placement of any storage object created during this operation.

| Name | Type | Description |
|---|---|---|
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| name | string | New name for consistency group. Required to resolve naming collisions. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot

A consistency group's snapshot

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group's snapshot to restore to. |
| uuid | string | The UUID of the consistency group's snapshot to restore to. |

restore_to

Use to restore a consistency group to a previous snapshot

| Name | Type | Description |
| --- | --- | --- |
| snapshot | snapshot | A consistency group's snapshot |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

space

Space information for the consistency group.

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in the consistency group, in bytes. |
| size | integer | The total provisioned size of the consistency group, in bytes. |
| used | integer | The amount of space consumed in the consistency group, in bytes. |

svm

The Storage Virtual Machine (SVM) in which the consistency group is located.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

object_stores

tiering

The tiering placement and policy definitions for volumes in this consistency group.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |
| win_sid_unix_id | string | Windows SID/UNIX ID depending on access-control type. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| shares | array[consistency_group_cifs_share] | |

export_clients

| Name | Type | Description |
|---|---|---|
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |

| Name | Type | Description |
| --- | --- | --- |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

export_policy

The policy associated with volumes to export them for protocol access.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| id | integer | Identifier for the export policy. |
| name | string | Name of the export policy. |
| rules | array[export_rules] | The set of rules that govern the export policy. |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

The CIFS share policy and/or export policies for this volume.

| Name | Type | Description |
| --- | --- | --- |
| cifs | cifs | |
| export_policy | export_policy | The policy associated with volumes to export them for protocol access. |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within an SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ UNIX-style security. |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number, consisting of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). First digit selects the set user ID (4), set group ID (2), and sticky (1) attributes. Second digit selects permission for the owner of the file. Third selects permissions for other users in the same group while the fourth selects permissions for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group, and other permissions are given (as in 755, representing the second, third and fourth digit), the first digit is assumed to be zero. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

qos

The QoS policy for this volume.

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available space, in bytes. |
| size | integer | Total provisioned size, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |

tiering

The tiering placement and policy definitions for this volume.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment for the volume. Valid in POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | The CIFS share policy and/or export policies for this volume. |
| qos | qos | The QoS policy for this volume. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | |
| tiering | tiering | The tiering placement and policy definitions for this volume. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |

consistency_groups

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |

| Name | Type | Description |
| --- | --- | --- |
| application | application | |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |
| name | string | Name of the consistency group. The consistency group name must be unique within an SVM.<br><br>If not provided and the consistency group contains only one volume, the name will be generated based on the volume name. If the consistency group contains more than one volume, the name is required. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| qos | qos | |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | Space information for the consistency group. |
| svm | svm | The Storage Virtual Machine (SVM) in which the consistency group is located. |
| tiering | tiering | The tiering placement and policy definitions for volumes in this consistency group. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency observed at the storage object, in microseconds.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| available_space | integer | The total space available in the consistency group, in bytes. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency observed at the storage object, in microseconds. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance and capacity data. |
| used_space | integer | The total space used in the consistency group, in bytes. |

replication_relationships

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| is_protected_by_svm_dr | boolean | Indicates whether or not this consistency group is protected by SVM DR. |
| is_source | boolean | Indicates whether or not this consistency group is the source for replication. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SnapMirror relationship. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| available_space | integer | The total space available in the consistency group, in bytes. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |
| used_space | integer | The total used space in the consistency group, in bytes. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a consistency group

PATCH `/application/consistency-groups/{uuid}`

**Introduced In:** 9.10

Updates a consistency group.

> that this operation will never delete storage elements. You can specify only elements that should be added to the consistency group regardless of existing storage objects.
> Mapping or unmapping a consistency group from igroups or subsystems is not supported.

#### Related ONTAP commands

- `vserver consistency-group modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the consistency group to modify. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |

| Name | Type | Description |
| --- | --- | --- |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group. The maximum number of items for this array is 16. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| vdisk_type | string | Type of objects in the consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

**Example request**


```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "application": {
    "component_type": "string",
    "type": "string"
  },
  "clone": {
    "parent_svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "consistency_groups": [
    {
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "application": {
        "component_type": "string",
        "type": "string"
      },
      "luns": [
        {
          "clone": {
            "source": {
              "name": "/vol/volume1/lun1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "comment": "string",
          "create_time": "2018-06-04 19:00:00 +0000",
          "lun_maps": [
            {
              "igroup": {
                "comment": "string",
                "igroups": [
                  {
                    "name": "igroup1",
```

```
                "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
              }
            ],
            "initiators": [
              {
                "comment": "my comment",
                "name": "iqn.1998-01.com.corp.iscsi:name1"
              }
            ],
            "name": "igroup1",
            "os_type": "string",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "provisioning_options": {
        "action": "string"
      },
      "qos": {
        "policy": {
          "max_throughput_iops": 10000,
          "max_throughput_mbps": 500,
          "min_throughput_iops": 2000,
          "min_throughput_mbps": 500,
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "serial_number": "string",
      "space": {
        "size": 1073741824,
        "used": 0
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "namespaces": [
    {
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "name": "/vol/volume1/qtree1/namespace1",
      "provisioning_options": {
        "action": "string"
      },
      "space": {
        "size": 1073741824,
```

```
            "used": 0
          },
          "subsystem_map": {
            "anagrpid": "00103050h",
            "nsid": "00000001h",
            "subsystem": {
              "comment": "string",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "parent_consistency_group": {
        "name": "my_consistency_group",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "provisioning_options": {
        "action": "string",
        "name": "string",
        "storage_service": {
          "name": "string"
        }
      },
      "qos": {
        "policy": {
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "restore_to": {
        "snapshot": {
          "name": "string",
          "uuid": "string"
        }
      },
      "snapshot_policy": {
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volumes": [
        {
          "comment": "string",
          "name": "vol_cs_dept",
          "nas": {
```

```
            "cifs": {
              "shares": [
                {
                  "comment": "HR Department Share",
                  "dir_umask": 18,
                  "file_umask": 18,
                  "home_directory": true,
                  "offline_files": "string",
                  "unix_symlink": "string",
                  "vscan_profile": "string"
                }
              ]
            },
            "export_policy": {
              "id": 0,
              "name": "string",
              "rules": [
                {
                  "anonymous_user": "string",
                  "chown_mode": "string",
                  "clients": [
                    {
                      "match": "0.0.0.0/0"
                    }
                  ],
                  "ntfs_unix_security": "string",
                  "protocols": [
                    "string"
                  ],
                  "ro_rule": [
                    "string"
                  ],
                  "rw_rule": [
                    "string"
                  ],
                  "superuser": [
                    "string"
                  ]
                }
              ]
            }
          },
          "provisioning_options": {
            "action": "string",
            "storage_service": {
              "name": "string"
```

```
            }
          },
          "qos": {
            "policy": {
              "name": "performance",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          },
          "snapshot_policy": {
            "name": "default",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "space": {
            "available": 0,
            "used": 0
          },
          "tiering": {
            "control": "string",
            "object_stores": [
              {
                "name": "string"
              }
            ]
          },
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ]
    }
  ],
  "luns": [
    {
      "clone": {
        "source": {
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "lun_maps": [
        {
          "igroup": {
            "comment": "string",
            "igroups": [
              {
                "name": "igroup1",
```

```
                "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
              }
            ],
            "initiators": [
              {
                "comment": "my comment",
                "name": "iqn.1998-01.com.corp.iscsi:name1"
              }
            ],
            "name": "igroup1",
            "os_type": "string",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "provisioning_options": {
        "action": "string"
      },
      "qos": {
        "policy": {
          "max_throughput_iops": 10000,
          "max_throughput_mbps": 500,
          "min_throughput_iops": 2000,
          "min_throughput_mbps": 500,
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "serial_number": "string",
      "space": {
        "size": 1073741824,
        "used": 0
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "namespaces": [
    {
      "comment": "string",
      "create_time": "2018-06-04 19:00:00 +0000",
      "name": "/vol/volume1/qtree1/namespace1",
      "provisioning_options": {
        "action": "string"
      },
      "space": {
        "size": 1073741824,
```

```
        "used": 0
      },
      "subsystem_map": {
        "anagrpid": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "comment": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "parent_consistency_group": {
    "name": "my_consistency_group",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "provisioning_options": {
    "action": "string",
    "name": "string",
    "storage_service": {
      "name": "string"
    }
  },
  "qos": {
    "policy": {
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "replication_relationships": [
    {
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "restore_to": {
    "snapshot": {
      "name": "string",
      "uuid": "string"
    }
  },
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
```

```
  "vdisk_type": "luns",
  "volumes": [
    {
      "comment": "string",
      "name": "vol_cs_dept",
      "nas": {
        "cifs": {
          "shares": [
            {
              "comment": "HR Department Share",
              "dir_umask": 18,
              "file_umask": 18,
              "home_directory": true,
              "offline_files": "string",
              "unix_symlink": "string",
              "vscan_profile": "string"
            }
          ]
        },
        "export_policy": {
          "id": 0,
          "name": "string",
          "rules": [
            {
              "anonymous_user": "string",
              "chown_mode": "string",
              "clients": [
                {
                  "match": "0.0.0.0/0"
                }
              ],
              "ntfs_unix_security": "string",
              "protocols": [
                "string"
              ],
              "ro_rule": [
                "string"
              ],
              "rw_rule": [
                "string"
              ],
              "superuser": [
                "string"
              ]
            }
          ]
```

```
        }
      },
      "provisioning_options": {
        "action": "string",
        "storage_service": {
          "name": "string"
        }
      },
      "qos": {
        "policy": {
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "snapshot_policy": {
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "space": {
        "available": 0,
        "used": 0
      },
      "tiering": {
        "control": "string",
        "object_stores": [
          {
            "name": "string"
          }
        ]
      },
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262285 | Consistency group does not support removing elements using a PATCH request. |
| 2621761 | Consistency groups with DP volumes are not supported on storage-limit enabled SVM. |
| 53411842 | Consistency group does not exist. |
| 53411843 | A consistency group with specified UUID was not found. |
| 53411844 | Specified consistency group was not found in the specified SVM. |
| 53411845 | The specified UUID and name refer to different consistency groups. |
| 53411846 | Either name or UUID must be provided. |
| 53411852 | A consistency group with the same identifier in the same scope exists. |
| 53411853 | Fields provided in the request conflict with each other. |
| 53411856 | Field provided is only supported when provisioning new objects. |
| 53411857 | LUNs that are not members of the application are not supported by this API. LUNs can be added to an application by adding the volume containing the LUNs to the application. |
| 53411860 | An object with the same identifier in the same scope exists. |
| 53411861 | The specified volume does not exist in provided volume array. |
| 53411862 | Modifying existing igroups is not supported using this API. |
| 53411864 | Request content insufficient to add an existing volume to an application. |
| 53411865 | Volumes contained in one consistency group cannot be added to a different consistency group. |
| 53411866 | LUNs are not supported on FlexGroup volumes. |
| 53411867 | LUN name is too long after appending a unique suffix. |

| Error Code | Description |
| --- | --- |
| 53411869 | Volume name is too long after appending a unique suffix. |
| 53411870 | When using the "round_robin" layout, the volume count must not be greater than the LUN count. |
| 53411942 | The application or component type of a consistency group that has an associated SnapMirror relationship cannot be changed. |
| 53411959 | Volumes with snapshot locking enabled cannot be added to a consistency group. |
| 53412027 | Failed to update the snapshot policy because the snapshot policies are not supported on the destination consistency group of SnapMirror active sync relationships. |
| 53412056 | The consistency group is not a FlexClone. |
| 53412057 | Consistency group split operation failed. |
| 53412058 | Failed to map or unmap because the consistency group consists of both 'luns' and 'namespaces'. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

application

| Name | Type | Description |
|---|---|---|
| component_type | string | Nested consistency group tag. |
| type | string | Top level consistency group tag. |

guarantee

| Name | Type | Description |
|---|---|---|
| type | string | The type of space guarantee of this volume in the aggregate. |

parent_consistency_group

Consistency group that is to be cloned.

_links

parent_snapshot

Consistency group that is to be cloned.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume name suffix/prefix for the cloned volumes.

| Name | Type | Description |
|---|---|---|
| prefix | string | Volume name prefix for cloned volumes. |
| suffix | string | Volume name suffix for cloned volumes. |

clone

Creates a clone of an existing consistency group from the current contents or an existing snapshot.

| Name | Type | Description |
|---|---|---|
| is_flexclone | boolean | Specifies if this consistency group is a FlexClone of a consistency group. |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| split_complete_percent | integer | Percentage of FlexClone blocks split from its parent consistency group. |
| split_estimate | integer | Space required to split the FlexClone consistency group. |
| split_initiated | boolean | Splits volumes after cloning. Defaults to false during POST. Only accepts true during a PATCH. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |
| uuid | string | The unique identifier of the clone source LUN.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

| Name | Type | Description |
|---|---|---|
| source | source | The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.<br><br>Valid in POST to create a new LUN as a clone of the source.<br><br>Valid in PATCH to overwrite an existing LUN's data as a clone of another. |

igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
|---|---|---|
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |

policy

The QoS policy

| Name | Type | Description |
|---|---|---|
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH. |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not reduced using the REST interface.<br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform, and the available space in the containing volume and aggregate.<br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

luns

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the `state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| lun_maps | array[lun_maps] | An array of LUN maps.<br><br>A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br><br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br><br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not reduced using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |

| Name | Type | Description |
| --- | --- | --- |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br><br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

consistency_group_nvme_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |

consistency_group_nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`.
They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | consistency_group_nvme_subsystem | An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts. |

namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing

the namespace.

| Name | Type | Description |
|---|---|---|
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| space | space | The storage space related properties of the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`.<br>They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. |
| uuid | string | The unique identifier of the NVMe namespace. |

parent_consistency_group

The parent consistency group.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the consistency group. |

storage_service

Determines the placement of any storage object created during this operation.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |
| name | string | New name for consistency group. Required to resolve naming collisions. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

policy

The QoS policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot

A consistency group's snapshot

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group's snapshot to restore to. |
| uuid | string | The UUID of the consistency group's snapshot to restore to. |

restore_to

Use to restore a consistency group to a previous snapshot

| Name | Type | Description |
| --- | --- | --- |
| snapshot | snapshot | A consistency group's snapshot |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

space

Space information for the consistency group.

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in the consistency group, in bytes. |
| size | integer | The total provisioned size of the consistency group, in bytes. |
| used | integer | The amount of space consumed in the consistency group, in bytes. |

svm

The Storage Virtual Machine (SVM) in which the consistency group is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for volumes in this consistency group.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
|---|---|---|
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |

| Name | Type | Description |
| --- | --- | --- |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| shares | array[consistency_group_cifs_share] | |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

export_policy

The policy associated with volumes to export them for protocol access.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Identifier for the export policy. |
| name | string | Name of the export policy. |
| rules | array[export_rules] | The set of rules that govern the export policy. |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the parent volume. |

nas

The CIFS share policy and/or export policies for this volume.

| Name | Type | Description |
| --- | --- | --- |
| cifs | cifs | |
| export_policy | export_policy | The policy associated with volumes to export them for protocol access. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| action | string | Operation to perform |
| count | integer | Number of elements to perform the operation on. |
| storage_service | storage_service | Determines the placement of any storage object created during this operation. |

qos

The QoS policy for this volume.

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The available space, in bytes. |
| size | integer | Total provisioned size, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |

tiering

The tiering placement and policy definitions for this volume.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | The CIFS share policy and/or export policies for this volume. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | The QoS policy for this volume. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| tiering | tiering | The tiering placement and policy definitions for this volume. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |

consistency_groups

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency observed at the storage object, in microseconds.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space.

| Name | Type | Description |
| --- | --- | --- |
| available_space | integer | The total space available in the consistency group, in bytes. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency observed at the storage object, in microseconds. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance and capacity data. |
| used_space | integer | The total space used in the consistency group, in bytes. |

replication_relationships

| Name | Type | Description |
| --- | --- | --- |
| is_protected_by_svm_dr | boolean | Indicates whether or not this consistency group is protected by SVM DR. |
| is_source | boolean | Indicates whether or not this consistency group is the source for replication. |
| uuid | string | The unique identifier of the SnapMirror relationship. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance and space numbers, such as, IOPS, latency, throughput, used space, and available space. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| available_space | integer | The total space available in the consistency group, in bytes. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency observed at the storage object, in microseconds. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| size | integer | The total size of the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |
| used_space | integer | The total used space in the consistency group, in bytes. |

consistency_group

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| application | application | |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | Creates a clone of an existing consistency group from the current contents or an existing snapshot. |
| consistency_groups | array[consistency_groups] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A consistency group can only be associated with one direct parent consistency group. |
| luns | array[luns] | The LUNs array can be used to create or modify LUNs in a consistency group on a new or existing volume that is a member of the consistency group. LUNs are considered members of a consistency group if they are located on a volume that is a member of the consistency group.<br>The maximum number of items for this array is 16. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the SVM using the NVMe over Fabrics protocol.<br>In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.<br><br>An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.<br>The maximum number of items for this array is 16.<br><br>• minItems: 0<br>• uniqueItems: 1<br>• Introduced in: 9.10 |
| parent_consistency_group | parent_consistency_group | The parent consistency group. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos | qos | |

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | Indicates whether or not replication has been enabled on this consistency group. |
| replication_relationships | array[replication_relationships] | Indicates the SnapMirror relationship of this consistency group. |
| replication_source | boolean | Since support for this field is to be removed in the next release, use replication_relationships.is_source instead. |
| restore_to | restore_to | Use to restore a consistency group to a previous snapshot |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| uuid | string | The unique identifier of the consistency group. The UUID is generated by ONTAP when the consistency group is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| vdisk_type | string | Type of objects in the consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| volumes | array[volumes] | A consistency group is a mutually exclusive aggregation of volumes or other consistency groups. A volume can only be associated with one direct parent consistency group.<br><br>The volumes array can be used to create new volumes in the consistency group, add existing volumes to the consistency group, or modify existing volumes that are already members of the consistency group.<br><br>The total number of volumes across all child consistency groups contained in a consistency group is constrained by the same limit. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve consistency group metrics

### Retrieve consistency group metrics

#### Overview

Retrieves historical performance and capacity metrics for a consistency group.

#### Examples

### Retrieving historical performance and space metrics for a consistency group for the past day

```
# The API:
/api/application/consistency-groups/{uuid}/metrics

# The call:
curl -X GET "https://<mgmt-ip>/api/application/consistency-
groups/{uuid}/metrics?interval=1d&fields=**" -H "accept:
application/hal+json"

# The response:
Snipped most of the records
{
"records": [
  {
    "uuid": "5069acd5-b325-11ed-a958-005056ac6b54",
    "timestamp": "2023-02-23T03:17:45Z",
    "status": "ok",
    "duration": "PT15S",
    "latency": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "iops": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "throughput": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "available_space": 862216192,
    "used_space": 1810432,
    "size": 864026624,
    "_links": {
```

```
      "self": {
        "href": "/api/application/consistency-groups/5069acd5-b325-11ed-
a958-005056ac6b54/metrics/2023-02-23T03%3A17%3A45Z?fields=**"
      }
    }
  },
  {
    "uuid": "5069acd5-b325-11ed-a958-005056ac6b54",
    "timestamp": "2023-02-23T02:18:00Z",
    "svm": {
      "name": "vs0",
      "uuid": "55b1a7b9-b2fb-11ed-a958-005056ac6b54"
    },
    "status": "partial_no_data",
    "duration": "PT15S",
    "latency": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "iops": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "throughput": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "available_space": 0,
    "used_space": 0,
    "size": 0,
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/5069acd5-b325-11ed-
a958-005056ac6b54/metrics/2023-02-23T02%3A18%3A00Z?fields=**"
      }
    }
  }
],
"num_records": 240,
"_links": {
```

```
  "self": {
    "href": "/api/application/consistency-groups/5069acd5-b325-11ed-a958-
005056ac6b54/metrics?fields=**"
  }
}
}
```

#### Retrieving historical metrics within a time range defined by timestamp

```
curl -X GET "https://<mgmt-ip>/api/application/consistency-
groups/{uuid}/metrics?fields=**&timestamp=2023-02-23T03:36:45Z..2023-02-
23T03:36:30Z" -H "accept: application/hal+json"

#### Response:
{
"records": [
{
  "uuid": "5069acd5-b325-11ed-a958-005056ac6b54",
  "timestamp": "2023-02-23T03:36:45Z",
  "svm": {
    "name": "vs0",
    "uuid": "55b1a7b9-b2fb-11ed-a958-005056ac6b54"
  },
  "status": "ok",
  "duration": "PT15S",
  "latency": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "available_space": 862183424,
  "used_space": 1843200,
```

```
  "size": 864026624,
  "_links": {
    "self": {
      "href": "/api/application/consistency-groups/5069acd5-b325-11ed-
a958-005056ac6b54/metrics/2023-02-23T03%3A36%3A45Z?fields=**"
    }
  }
},
{
  "uuid": "5069acd5-b325-11ed-a958-005056ac6b54",
  "timestamp": "2023-02-23T03:36:30Z",
  "svm": {
    "name": "vs0",
    "uuid": "55b1a7b9-b2fb-11ed-a958-005056ac6b54"
  },
  "status": "ok",
  "duration": "PT15S",
  "latency": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "available_space": 862183424,
  "used_space": 1843200,
  "size": 864026624,
  "_links": {
    "self": {
      "href": "/api/application/consistency-groups/5069acd5-b325-11ed-
a958-005056ac6b54/metrics/2023-02-23T03%3A36%3A30Z?fields=**"
    }
  }
}
],
```

```
"num_records": 2,
"_links": {
"self": {
  "href": "/api/application/consistency-groups/5069acd5-b325-11ed-a958-
005056ac6b54/metrics?fields=**&timestamp=2023-02-23T03:36:45Z..2023-02-
23T03:36:30Z"
}
}
}
```


### Retrieve the historical performance and capacity metrics for a consistency group

GET `/application/consistency-groups/{consistency_group.uuid}/metrics`

**Introduced In:** 9.13

Retrieves historical performance and capacity metrics for a consistency group.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| size | integer | query | False | Filter by size |
| duration | string | query | False | Filter by duration |
| available_space | integer | query | False | Filter by available_space |
| status | string | query | False | Filter by status |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.read | integer | query | False | Filter by iops.read |
| throughput.other | integer | query | False | Filter by throughput.other |
| throughput.total | integer | query | False | Filter by throughput.total |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| timestamp | string | query | False | Filter by timestamp |
| used_space | integer | query | False | Filter by used_space |
| consistency_group.uuid | string | path | True | Unique identifier of the consistency group. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1w, 1m, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "available_space": 4096,
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "size": 4096,
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "used_space": 4096
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency observed at the storage object, in microseconds.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance and capacity numbers, such as, IOPS, latency, throughput, used space, and available space.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| available_space | integer | The total space available in the consistency group, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency observed at the storage object, in microseconds. |
| size | integer | The total size of the consistency group, in bytes. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance and capacity data. |

| Name | Type | Description |
|---|---|---|
| used_space | integer | The total space used in the consistency group, in bytes. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage application consistency group snapshots

### Manage application consistency group snapshots

#### Overview

Consistency groups support snapshot create, inventory, and restore. Snapshots can be created on a specified schedule or on-demand. On-demand snapshots can have a type of application consistent or crash consistent. Crash consistent is the default. Scheduled snapshotss are always crash consistent. There is no functional difference in ONTAP between crash consistent or application consistent snapshots.

The functionality provided by these APIs is not integrated with the host application. Snapshots have limited value without host coordination, so the use of the SnapCenter Backup Management suite is recommended to ensure correct interaction between host applications and ONTAP.

### On-Demand Snapshots

A manual snapshot may be created on-demand for a parent consistency group and for any of the children consistency groups within it.

Scheduled and manual snapshot creation operations are subject to a pre-defined seven second internal timeout. If the snapshot creation operation does not complete within this time, it is aborted.

Individual volume snapshots within a consistency group snapshots can be accessed and used with native volume snapshot operations.

When an individual volume snapshot is deleted that is part of a consistency group snapshot, then that consistency group snapshot becomes invalid and which cannot be used for restoring the consistency group.

### Restoring to a Previous Snapshot

A snapshot restores to a parent consistency group from an existing parent consistency group's snapshot. A snapshot restores to any of the children's consistency groups within it from an existing children's consistency group. Granular snapshots are supported. This is performed by a PATCH operation on the specific consistency group for the restore. An example is shown in PATCH /application/consistency-groups .

Any existing snapshots that were created chronologically after the time of the snapshot used in a successful restore operation is deleted, in compliance with existing ONTAP "future-snapshot" handling principles.

On failures during consistency group restores, any volumes that have been restored will remain so and will not be rolled back. The user must retry the failed restore operation until it is successful. The user can retry with consistency group restore or individual volume-granular restore.

#### Consistency group Snapshot APIs

The following APIs are used to perform operations related to consistency group snapshots:

– GET /api/application/consistency-groups/{consistency_group.uuid}/snapshots

– POST /api/application/consistency-groups/{consistency_group.uuid}/snapshots

– POST /api/application/consistency-groups/{consistency_group.uuid}/snapshots?action=start

– GET /api/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}

– PATCH /api/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}?action=commit

– DELETE /api/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}

#### Examples

### Required properties

- `consistency_group.uuid` - Existing consistency group UUID in which to create the snapshot.

### Retrieving the list of existing snapshots for a consistency group

Retrieves the list of consistency group granular snapshots for a specific consistency group.

```
curl -X GET 'https://<mgmt-ip>/api/application/consistency-
groups/92c6c770-17a1-11eb-b141-005056acd498/snapshots'

#### Response:
{
"records": [
  {
    "uuid": "92c6c770-17a1-11eb-b141-005056acd498",
    "name": "sa3s1",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/a8d0626a-17a0-11eb-
b141-005056acd498/snapshots/92c6c770-17a1-11eb-b141-005056acd498"
      }
    }
  },
  {
    "uuid": "c5a250ba-17a1-11eb-b141-005056acd498",
    "name": "sa3s2",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/a8d0626a-17a0-11eb-
b141-005056acd498/snapshots/c5a250ba-17a1-11eb-b141-005056acd498"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/application/consistency-groups/a8d0626a-17a0-11eb-b141-
005056acd498/snapshots"
  }
}
}
```

#### Retrieves details of a specific snapshot for a consistency group

Retrieves details for a specific snapshot in a consistency group.

```
curl -X GET 'https://<mgmt-ip>/api/application/consistency-
groups/92c6c770-17a1-11eb-b141-005056acd498/snapshots/a175c021-4199-11ec-
8674-005056accf3f'

#### Response:
{
"consistency_group": {
  "uuid": "ddabc6a5-4196-11ec-8674-005056accf3f",
  "name": "CG_1",
  "_links": {
    "self": {
      "href": "/api/application/consistency-groups/ddabc6a5-4196-11ec-
8674-005056accf3f"
    }
  }
},
"uuid": "a175c021-4199-11ec-8674-005056accf3f",
"name": "sa3s2",
"consistency_type": "crash",
"comment": "manually created snapshot",
"create_time": "2021-11-09T15:14:23-05:00",
"svm": {
  "uuid": "7379fecb-4195-11ec-8674-005056accf3f",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/7379fecb-4195-11ec-8674-005056accf3f"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/application/consistency-groups/ddabc6a5-4196-11ec-8674-
005056accf3f/snapshots/a175c021-4199-11ec-8674-005056accf3f"
  }
}
}
```

#### Retrieving bulk snapshots

Retrieves the list of consistency group granular snapshots for all consistency groups on the cluster.

```
curl -X GET 'https://<mgmt-ip>/api/application/consistency-
groups/*/snapshots'
```

```
#### Response:
{
"records": [
  {
    "consistency_group": {
      "uuid": "1218f900-c124-11ee-bbfe-005056acb65e",
      "name": "cg3",
      "_links": {
        "self": {
          "href": "/api/application/consistency-groups/1218f900-c124-11ee-
bbfe-005056acb65e"
        }
      }
    },
    "uuid": "7da4d364-c12e-11ee-bbfe-005056acb65e",
    "name": "cg3ss",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/1218f900-c124-11ee-
bbfe-005056acb65e/snapshots/7da4d364-c12e-11ee-bbfe-005056acb65e"
      }
    }
  },
  {
    "consistency_group": {
      "uuid": "15a8f66e-c124-11ee-bbfe-005056acb65e",
      "name": "cg2",
      "_links": {
        "self": {
          "href": "/api/application/consistency-groups/15a8f66e-c124-11ee-
bbfe-005056acb65e"
        }
      }
    },
    "uuid": "83595384-c12e-11ee-bbfe-005056acb65e",
    "name": "cg2ss",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/15a8f66e-c124-11ee-
bbfe-005056acb65e/snapshots/83595384-c12e-11ee-bbfe-005056acb65e"
      }
    }
  },
  {
    "consistency_group": {
      "uuid": "1c101d17-c124-11ee-bbfe-005056acb65e",
```

```
      "name": "cg1",
      "_links": {
        "self": {
          "href": "/api/application/consistency-groups/1c101d17-c124-11ee-
bbfe-005056acb65e"
        }
      }
    },
    "uuid": "87d0e49c-c12e-11ee-bbfe-005056acb65e",
    "name": "cg1ss",
    "_links": {
      "self": {
        "href": "/api/application/consistency-groups/1c101d17-c124-11ee-
bbfe-005056acb65e/snapshots/87d0e49c-c12e-11ee-bbfe-005056acb65e"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/application/consistency-groups/*/snapshots"
  }
}
}
```


#### Creating a crash-consistent snapshot of a consistency group

Creates an on-demand crash-consistent snapshot of an existing consistency group.

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-
groups/a8d0626a-17a0-11eb-b141-005056acd498/snapshots' -d '{ "name":
"name_of_this_snapshot", "consistency_type": "crash", "comment": "this is
a manually created on-demand snapshot", "snapmirror_label":
"my_special_sm_label" }' -H "accept: application/hal+json"

#### Response:
{
}
```

#### Creating a app-consistent snapshot of a consistency group

Creates an on-demand crash-consistent snapshot of an existing consistency group.

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-
groups/a8d0626a-17a0-11eb-b141-005056acd498/snapshots' -d '{ "name":
"name_of_this_snapshot", "consistency_type": "application", "comment":
"this is a manually created on-demand snapshot", "snapmirror_label":
"my_special_sm_label" }' -H "accept: application/hal+json"

#### Response:
{
}
```

#### Starting a two-phase crash-consistent snapshot creation for a consistency group

Starts a two-phase on-demand crash-consistent snapshot creation for an existing consistency group.

```
curl -X POST 'https://<mgmt-ip>/api/application/consistency-
groups/a8d0626a-17a0-11eb-b141-
005056acd498/snapshots?action=start&action_timeout=7' -d '{ "name":
"name_of_this_snapshot", "consistency_type": "application", "comment":
"this is a manually created on-demand snapshot", "snapmirror_label":
"my_special_sm_label" }' -H "accept: application/hal+json"

#### Response:
Location: /api/application/consistency-groups/a8d0626a-17a0-11eb-b141-
005056acd498/snapshots/7aac0607-0c4d-11ee-ad32-005056a73101
{
}
```

#### Committing a previously started two-phase crash-consistent snapshot creation for a consistency group

Commits a previously started two-phase on-demand crash-consistent snapshot creation for an existing consistency group.

```
curl -X PATCH 'https://<mgmt-ip>/api/application/consistency-
groups/a8d0626a-17a0-11eb-b141-005056acd498/snapshots/7aac0607-0c4d-11ee-
ad32-005056a73101?action=commit' -H "accept: application/hal+json"

#### Response:
{
}
```

#### Deleting a snapshot from a consistency group

Deletes an existing snapshot from a consistency group.

```
curl -X DELETE 'https://<mgmt-ip>/api/application/consistency-
groups/a8d0626a-17a0-11eb-b141-005056acd498/snapshots/92c6c770-17a1-11eb-
b141-005056acd498'

#### Response:
{
}
```

### Retrieve consistency group snapshots

GET `/application/consistency-groups/{consistency_group.uuid}/snapshots`

**Introduced In:** 9.10

Retrieves snapshots for a consistency group.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See DOC Requesting specific fields to learn more.

- `is_partial`
- `missing_volumes.uuid`
- `missing_volumes.name`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | path | True | The unique identifier of the consistency group to retrieve. |
| write_fence | boolean | query | False | Filter by write_fence<br>• Introduced in: 9.14 |
| snapmirror_label | string | query | False | Filter by snapmirror_label |
| name | string | query | False | Filter by name |
| missing_volumes.name | string | query | False | Filter by missing_volumes.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| missing_volumes.uuid | string | query | False | Filter by missing_volumes.uuid |
| luns.uuid | string | query | False | Filter by luns.uuid<br>• Introduced in: 9.16 |
| luns.name | string | query | False | Filter by luns.name<br>• Introduced in: 9.16 |
| missing_luns.uuid | string | query | False | Filter by missing_luns.uuid<br>• Introduced in: 9.16 |
| missing_luns.name | string | query | False | Filter by missing_luns.name<br>• Introduced in: 9.16 |
| consistency_type | string | query | False | Filter by consistency_type |
| uuid | string | query | False | Filter by uuid |
| comment | string | query | False | Filter by comment |
| missing_namespaces.uuid | string | query | False | Filter by missing_namespaces.uuid<br>• Introduced in: 9.16 |
| missing_namespaces.name | string | query | False | Filter by missing_namespaces.name<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| snapshot_volumes.snapshot.name | string | query | False | Filter by snapshot_volumes.snapshot.name<br><br>• Introduced in: 9.12 |
| snapshot_volumes.snapshot.uuid | string | query | False | Filter by snapshot_volumes.snapshot.uuid<br><br>• Introduced in: 9.12 |
| snapshot_volumes.volume.name | string | query | False | Filter by snapshot_volumes.volume.name<br><br>• Introduced in: 9.12 |
| snapshot_volumes.volume.uuid | string | query | False | Filter by snapshot_volumes.volume.uuid<br><br>• Introduced in: 9.12 |
| restore_size | integer | query | False | Filter by restore_size<br><br>• Introduced in: 9.16 |
| is_partial | boolean | query | False | Filter by is_partial |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| namespaces.uuid | string | query | False | Filter by namespaces.uuid<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespaces.name | string | query | False | Filter by namespaces.name<br>• Introduced in: 9.16 |
| consistency_group.name | string | query | False | Filter by consistency_group.name |
| reclaimable_space | integer | query | False | Filter by reclaimable_space<br>• Introduced in: 9.16 |
| create_time | string | query | False | Filter by create_time |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records. |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "My snapshot comment",
      "consistency_group": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "my_consistency_group",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "consistency_type": "crash",
      "create_time": "2020-10-25 11:20:00 +0000",
      "is_partial": "",
      "luns": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "missing_luns": [
        {
          "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "missing_namespaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/namespace1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "missing_volumes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ],
      "name": "string",
      "namespaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/namespace1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "reclaimable_space": 0,
      "restore_size": 4096,
      "snapmirror_label": "sm_label",
      "snapshot_volumes": [
```

```
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "snapshot": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "this_snapshot",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

consistency_group

The consistency group of the snapshot.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

luns

A reference to a LUN.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

lun_reference

A reference to a LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

missing_namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

volume_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br><br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

consistency_group_volume_snapshot

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| snapshot | snapshot_reference | |
| volume | volume_reference | |

svm_reference

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

records

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| comment | string | Comment for the snapshot. |
| consistency_group | consistency_group | The consistency group of the snapshot. |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Time the snapshot copy was created |
| is_partial | boolean | Indicates whether the snapshot taken is partial or not. |
| luns | array[luns] | The list of LUNs in this snapshot. |
| missing_luns | array[lun_reference] | List of LUNs that are not in the snapshot. |
| missing_namespaces | array[missing_namespaces] | List of NVMe namespaces that are not in the snapshot. |
| missing_volumes | array[volume_reference] | List of volumes which are not in the snapshot. |
| name | string | Name of the snapshot. |
| namespaces | array[namespaces] | The list of NVMe namespaces in this snapshot. |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| restore_size | integer | Size of the consistency group if this snapshot is restored. |
| snapmirror_label | string | Snapmirror Label for the snapshot. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_volumes | array[consistency_group_volume_snapshot] | List of volume and snapshot identifiers for each volume in the snapshot. |
| svm | svm_reference | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the snapshot. The UUID is generated by ONTAP when the snapshot is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| write_fence | boolean | Specifies whether a write fence will be taken when creating the snapshot. The default is false if there is only one volume in the consistency group, otherwise the default is true. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a consistency group snapshot

POST `/application/consistency-groups/{consistency_group.uuid}/snapshots`

**Introduced In:** 9.10

Creates a snapshot of an existing consistency group.

#### Required properties

- `consistency_group.uuid` - Existing consistency group UUID in which to create the snapshot.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | path | True | The unique identifier of the consistency group to retrieve. |
| action | string | query | False | Initiates the snapshot create operation. The start of the snapshot operation can optionally use a timeout value specified by "action_timeout". The snapshot is committed by calling PATCH on the snapshot href link with action specified as "commit".<br>• Introduced in: 9.11<br>• enum: ["start"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| action_timeout | integer | query | False | Duration to complete the 2-phase snapshot operation. This also specifies the maximum duration that the write-fence remains in effect on the volumes associated with this consistency group. Default is 7 seconds with a valid range of 5 to 120 seconds.<br>• Introduced in: 9.11 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | Comment for the snapshot. |
| consistency_group | consistency_group | The consistency group of the snapshot. |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Time the snapshot copy was created |
| is_partial | boolean | Indicates whether the snapshot taken is partial or not. |
| luns | array[luns] | The list of LUNs in this snapshot. |
| missing_luns | array[lun_reference] | List of LUNs that are not in the snapshot. |
| missing_namespaces | array[missing_namespaces] | List of NVMe namespaces that are not in the snapshot. |
| missing_volumes | array[volume_reference] | List of volumes which are not in the snapshot. |
| name | string | Name of the snapshot. |
| namespaces | array[namespaces] | The list of NVMe namespaces in this snapshot. |

| Name | Type | Description |
| --- | --- | --- |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| restore_size | integer | Size of the consistency group if this snapshot is restored. |
| snapmirror_label | string | Snapmirror Label for the snapshot. |
| snapshot_volumes | array[consistency_group_volume_snapshot] | List of volume and snapshot identifiers for each volume in the snapshot. |
| svm | svm_reference | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the snapshot. The UUID is generated by ONTAP when the snapshot is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| write_fence | boolean | Specifies whether a write fence will be taken when creating the snapshot. The default is false if there is only one volume in the consistency group, otherwise the default is true. |

**Example request**

```
{
  "comment": "My snapshot comment",
  "consistency_group": {
    "name": "my_consistency_group",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "consistency_type": "crash",
  "create_time": "2020-10-25 11:20:00 +0000",
  "is_partial": "",
  "luns": [
    {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_luns": [
    {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_namespaces": [
    {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_volumes": [
    {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "name": "string",
  "namespaces": [
    {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "reclaimable_space": 0,
  "restore_size": 4096,
  "snapmirror_label": "sm_label",
  "snapshot_volumes": [
```

```
      {
        "snapshot": {
          "name": "this_snapshot",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response


```
{
  "job": {
    "uuid": "string"
  }
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53411918 | Snapshot operation not permitted. |
| 53411921 | Failed to create snapshot for consistency group. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

consistency_group

The consistency group of the snapshot.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

luns

A reference to a LUN.

| Name | Type | Description |
|---|---|---|
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

lun_reference

A reference to a LUN.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

missing_namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

volume_reference

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br><br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

consistency_group_volume_snapshot

| Name | Type | Description |
| --- | --- | --- |
| snapshot | snapshot_reference | |
| volume | volume_reference | |

svm_reference

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

consistency_group_snapshot

| Name | Type | Description |
|---|---|---|
| comment | string | Comment for the snapshot. |
| consistency_group | consistency_group | The consistency group of the snapshot. |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Time the snapshot copy was created |
| is_partial | boolean | Indicates whether the snapshot taken is partial or not. |
| luns | array[luns] | The list of LUNs in this snapshot. |
| missing_luns | array[lun_reference] | List of LUNs that are not in the snapshot. |
| missing_namespaces | array[missing_namespaces] | List of NVMe namespaces that are not in the snapshot. |
| missing_volumes | array[volume_reference] | List of volumes which are not in the snapshot. |
| name | string | Name of the snapshot. |

| Name | Type | Description |
| --- | --- | --- |
| namespaces | array[namespaces] | The list of NVMe namespaces in this snapshot. |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| restore_size | integer | Size of the consistency group if this snapshot is restored. |
| snapmirror_label | string | Snapmirror Label for the snapshot. |
| snapshot_volumes | array[consistency_group_volume_snapshot] | List of volume and snapshot identifiers for each volume in the snapshot. |
| svm | svm_reference | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the snapshot. The UUID is generated by ONTAP when the snapshot is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| write_fence | boolean | Specifies whether a write fence will be taken when creating the snapshot. The default is false if there is only one volume in the consistency group, otherwise the default is true. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a consistency group snapshot

DELETE `/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}`

**Introduced In:** 9.10

Deletes a snapshot of a consistency group.

#### Examples

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | path | True | The unique identifier of the snapshot of the consistency group to delete. |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53412007 | Failed to delete the snapshot because it has not expired or is locked. |
| 53412008 | Failed to delete the consistency group snapshot because it is currently used as a reference for a replication relationship. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a consistency group snapshot

GET `/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}`

**Introduced In:** 9.10

Retrieves details of a specific snapshot for a consistency group.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See DOC Requesting specific fields to learn more.

- `is_partial`
- `missing_volumes.uuid`
- `missing_volumes.name`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | path | True | The unique identifier of the consistency group to retrieve. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | The unique identifier of the snapshot of the consistency group to retrieve. |
| write_fence | boolean | query | False | Filter by write_fence<br>• Introduced in: 9.14 |
| snapmirror_label | string | query | False | Filter by snapmirror_label |
| name | string | query | False | Filter by name |
| missing_volumes.name | string | query | False | Filter by missing_volumes.name |
| missing_volumes.uuid | string | query | False | Filter by missing_volumes.uuid |
| luns.uuid | string | query | False | Filter by luns.uuid<br>• Introduced in: 9.16 |
| luns.name | string | query | False | Filter by luns.name<br>• Introduced in: 9.16 |
| missing_luns.uuid | string | query | False | Filter by missing_luns.uuid<br>• Introduced in: 9.16 |
| missing_luns.name | string | query | False | Filter by missing_luns.name<br>• Introduced in: 9.16 |
| consistency_type | string | query | False | Filter by consistency_type |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| comment | string | query | False | Filter by comment |
| missing_namespaces.uuid | string | query | False | Filter by missing_namespaces.uuid<br><br>• Introduced in: 9.16 |
| missing_namespaces.name | string | query | False | Filter by missing_namespaces.name<br><br>• Introduced in: 9.16 |
| snapshot_volumes.snapshot.name | string | query | False | Filter by snapshot_volumes.snapshot.name<br><br>• Introduced in: 9.12 |
| snapshot_volumes.snapshot.uuid | string | query | False | Filter by snapshot_volumes.snapshot.uuid<br><br>• Introduced in: 9.12 |
| snapshot_volumes.volume.name | string | query | False | Filter by snapshot_volumes.volume.name<br><br>• Introduced in: 9.12 |
| snapshot_volumes.volume.uuid | string | query | False | Filter by snapshot_volumes.volume.uuid<br><br>• Introduced in: 9.12 |
| restore_size | integer | query | False | Filter by restore_size<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| is_partial | boolean | query | False | Filter by is_partial |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| namespaces.uuid | string | query | False | Filter by namespaces.uuid<br>• Introduced in: 9.16 |
| namespaces.name | string | query | False | Filter by namespaces.name<br>• Introduced in: 9.16 |
| consistency_group.name | string | query | False | Filter by consistency_group.name |
| reclaimable_space | integer | query | False | Filter by reclaimable_space<br>• Introduced in: 9.16 |
| create_time | string | query | False | Filter by create_time |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| comment | string | Comment for the snapshot. |
| consistency_group | consistency_group | The consistency group of the snapshot. |
| consistency_type | string | Consistency type. This is for categorization purposes only. A snapshot should not be set to 'application consistent' unless the host application is quiesced for the snapshot. Valid in POST. |
| create_time | string | Time the snapshot copy was created |

| Name | Type | Description |
| --- | --- | --- |
| is_partial | boolean | Indicates whether the snapshot taken is partial or not. |
| luns | array[luns] | The list of LUNs in this snapshot. |
| missing_luns | array[lun_reference] | List of LUNs that are not in the snapshot. |
| missing_namespaces | array[missing_namespaces] | List of NVMe namespaces that are not in the snapshot. |
| missing_volumes | array[volume_reference] | List of volumes which are not in the snapshot. |
| name | string | Name of the snapshot. |
| namespaces | array[namespaces] | The list of NVMe namespaces in this snapshot. |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| restore_size | integer | Size of the consistency group if this snapshot is restored. |
| snapmirror_label | string | Snapmirror Label for the snapshot. |
| snapshot_volumes | array[consistency_group_volume_snapshot] | List of volume and snapshot identifiers for each volume in the snapshot. |
| svm | svm_reference | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the snapshot. The UUID is generated by ONTAP when the snapshot is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| write_fence | boolean | Specifies whether a write fence will be taken when creating the snapshot. The default is false if there is only one volume in the consistency group, otherwise the default is true. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "My snapshot comment",
  "consistency_group": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "my_consistency_group",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "consistency_type": "crash",
  "create_time": "2020-10-25 11:20:00 +0000",
  "is_partial": "",
  "luns": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_luns": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_namespaces": [
    {
      "_links": {
```

```
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "missing_volumes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "name": "string",
  "namespaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "reclaimable_space": 0,
  "restore_size": 4096,
  "snapmirror_label": "sm_label",
  "snapshot_volumes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "snapshot": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
          },
          "name": "this_snapshot",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

consistency_group

The consistency group of the snapshot.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

luns

A reference to a LUN.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the LUN. |

lun_reference

A reference to a LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

missing_namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

volume_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

namespaces

A reference to a namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of an NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the NVMe namespace. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

consistency_group_volume_snapshot

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| snapshot | snapshot_reference | |
| volume | volume_reference | |

svm_reference

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Complete a consistency group snapshot operation

PATCH `/application/consistency-groups/{consistency_group.uuid}/snapshots/{uuid}`

**Introduced In:** 9.11

Completes a snapshot operation of a consistency group.

#### Example

#### Completing a snapshot operation

The following example shows how to complete the snapshot operation by committing an existing snapshot to disk:

```
curl -X PATCH https://<mgmt-ip>/api/application/consistency-groups/a8d0626a-17a0-11eb-b141-005056acd498/snapshots/92c6c770-17a1-11eb-b141-005056acd498?action=commit
```

### Response:

```
{
}
```

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | path | True | Unique identifier of the consistency group's snapshot for deletion. |
| uuid | string | path | True | Snapshot UUID. |
| action | string | query | False | Commits the snapshot. The commit must be called within the timeout value specified during POST. If the commit is not called within the specified time, then the snapshot create operation is aborted.<br>• enum: ["commit"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53411925 | Failed to find a previously initiated two-phase snapshot operation for consistency group. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage application containers

### Manage application containers

#### Overview

Application containers provision one or more storage objects. Currently, only NAS volumes are supported. Application containers allow you to specify the policies and rules for enabling and managing client access to storage. FlexCache volumes can also be provisioned.

#### Examples

**Creating a FlexVol with NAS (NFS and CIFS access) along with S3 NAS bucket with S3 access policies**

```
# The API:
/api/application/containers

# The call:
curl -X POST 'https://<mgmt-ip>/api/application/containers' -d '{ "svm": {
"name": "vs0" }, "volumes": [ { "name": "vol1", "space": { "size": "100mb"
}, "scale_out": "false", "nas": { "path": "/vol1", "export_policy": {
"name": "vol1", "rules": [ { "clients": [ { "match": "0.0.0.0/0" } ],
"rw_rule": [ "any" ], "ro_rule": [ "any" ] } ] }, "cifs": { "shares": [ {
"name": "vol1", "acls": [ { "type": "windows", "permission":
"full_control", "user_or_group": "everyone" } ] } ] } }, "s3_bucket": {
"name": "vol1", "nas_path": "/vol1", "policy": { "statements": [ {
"actions": [ "ListBucket" ], "effect": "allow", "principals": [ "user1",
"group/grp1" ], "resources": [ "vol1", "vol1/*" ] } ] } } } ] }'

#### Response:
{
"job": {
  "uuid": "9c9cabf3-0a88-11ec-a449-005056bbcf9f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9c9cabf3-0a88-11ec-a449-005056bbcf9f"
    }
  }
}
}
```

**Create an application container**

POST `/application/containers`

**Introduced In:** 9.17

Creates one or more of the following:

- New NAS FlexVol or FlexGroup volumes
- S3 buckets
- Access policies for NFS, CIFS and S3
- FlexCache volumes

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the container.
- `volumes`

#### Naming Conventions

##### Volume

- volumes[].name, if specified
- suffixed by "_#" where "#" is a system generated unique number, if provisioning_options.count is provided

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| svm | svm | The SVM in which the container is located. |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning the volume. Only mirrored aggregates are used if this parameter is set to *true* and only unmirrored aggregates are used if this parameter is set to *false*. The default value is *true* for a MetroCluster configuration and is *false* for a non-MetroCluster configuration. |
| volumes | array[volumes] | A list of NAS volumes to provision. |

**Example request**

```
{
  "provisioning_options": {
    "exclude_aggregates": [
      {
        "name": "aggr1"
      }
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volumes": [
    {
      "exclude_aggregates": [
        {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "flexcache": {
        "origins": [
          {
            "svm": {
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "volume": {
              "name": "volume1",
              "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
            }
          }
        ]
      },
      "name": "vol_cs_dept",
      "nas": {
        "cifs": {
          "shares": [
            {
              "acls": [
                {
                  "permission": "string",
                  "type": "string",
                  "user_or_group": "ENGDOMAIN\\ad_user"
```

```
                }
              ],
              "comment": "HR Department Share",
              "dir_umask": 18,
              "file_umask": 18,
              "name": "HR_SHARE",
              "offline_files": "string",
              "unix_symlink": "string",
              "vscan_profile": "string"
            }
          ]
        },
        "export_policy": {
          "id": 0,
          "name": "string",
          "rules": [
            {
              "anonymous_user": "string",
              "chown_mode": "string",
              "clients": [
                {
                  "match": "0.0.0.0/0"
                }
              ],
              "ntfs_unix_security": "string",
              "protocols": [
                "string"
              ],
              "ro_rule": [
                "string"
              ],
              "rw_rule": [
                "string"
              ],
              "superuser": [
                "string"
              ]
            }
          ]
        },
        "junction_parent": {
          "name": "vs1_root",
          "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
        },
        "path": "/user/my_volume",
        "security_style": "string",
```

```
        "unix_permissions": 493
      },
      "qos": {
        "policy": {
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "s3_bucket": {
        "name": "bucket1",
        "nas_path": "/",
        "policy": {
          "statements": [
            {
              "actions": [
                "GetObject",
                "PutObject",
                "DeleteObject",
                "ListBucket"
              ],
              "conditions": [
                {
                  "delimiters": [
                    "/"
                  ],
                  "max_keys": [
                    1000
                  ],
                  "operator": "ip_address",
                  "prefixes": [
                    "pref"
                  ],
                  "source_ips": [
                    "1.1.1.1",
                    "1.2.2.0/24"
                  ],
                  "usernames": [
                    "user1"
                  ]
                }
              ],
              "effect": "allow",
              "principals": [
                "user1",
                "group/grp1",
                "nasgroup/group1"
```

```
                ],
                "resources": [
                  "bucket1",
                  "bucket1/*"
                ],
                "sid": "Full_Access_To_User1!"
              }
            ]
          }
        },
        "snaplock": {
          "append_mode_enabled": "",
          "autocommit_period": "P30M",
          "retention": {
            "default": "P30Y",
            "maximum": "P30Y",
            "minimum": "P30Y"
          },
          "type": "enterprise"
        },
        "snapshot_policy": {
          "name": "default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "space": {
          "size": 1073741824
        },
        "storage_service": {
          "name": "string"
        },
        "tiering": {
          "control": "string",
          "object_stores": [
            {
              "name": "string"
            }
          ],
          "policy": "string"
        }
      }
    ]
  }
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13109173 | Failed to access cluster capabilities. |
| 53412079 | Effective Cluster Version not supported for the request. |
| 53412080 | CIFS shares name already exists. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

exclude_aggregates

| Name | Type | Description |
|---|---|---|
| name | string | |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

The SVM in which the container is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

exclude_aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

container_volume_flexcache_relationship

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| volume | volume | |

writeback

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether or not writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume. |

flexcache

The FlexCache origin volume.

| Name | Type | Description |
|---|---|---|
| dr_cache | boolean | If set to true, a DR cache is created. |
| origins | array[container_volume_flexcache_relationship] | |
| writeback | writeback | |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

consistency_group_cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
|---|---|---|
| access_based_enumeration | boolean | Specifies whether all folders inside this share are visible to a user based on that individual user's access right; prevents the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.12<br>• x-nullable: true |
| dir_umask | integer | Directory mode creation mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | boolean | Specifies whether SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | integer | File mode creation mask to be viewed as an octal number. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.12<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow Unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specifies whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| shares | array[consistency_group_cifs_share] | |

self_link

export_clients

| Name | Type | Description |
|---|---|---|
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
|---|---|---|
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |

| Name | Type | Description |
| --- | --- | --- |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

export_policy

The policy associated with volumes to export them for protocol access.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Identifier for the export policy. |
| name | string | Name of the export policy. |
| rules | array[export_rules] | The set of rules that govern the export policy. |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

The CIFS share policy and export policies for this volume.

| Name | Type | Description |
| --- | --- | --- |
| cifs | cifs | |
| export_policy | export_policy | The policy associated with volumes to export them for protocol access. |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within an SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ UNIX-style security. |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number, consisting of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). First digit selects the set user ID (4), set group ID (2), and sticky (1) attributes. Second digit selects permission for the owner of the file. Third selects permissions for other users in the same group while the fourth selects permissions for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group, and other permissions are given (as in 755, representing the second, third and fourth digit), the first digit is assumed to be zero. |

policy

The QoS policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | The QoS policy |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |

| Name | Type | Description |
| --- | --- | --- |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

statements

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies are evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[statements] | Specifies the bucket access policy statement. |

s3_bucket

The S3 bucket

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path that corresponds with the NAS bucket. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies are evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num> |
| retention | retention | |
| type | string | The SnapLock type of the volume.<br> compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files.<br> enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume.<br> non_snaplock ‐ Indicates the volume is non-snaplock. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

space

| Name | Type | Description |
|---|---|---|
| size | integer | The total provisioned size of the container, in bytes. |

storage_service

Determines the placement of the volume that is to be provisioned.

| Name | Type | Description |
|---|---|---|
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

| Name | Type | Description |
|---|---|---|
| name | string | The name of the object store to use. Used for placement. |

tiering

| Name | Type | Description |
|---|---|---|
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |
| flexcache | flexcache | The FlexCache origin volume. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | The CIFS share policy and export policies for this volume. |
| qos | qos | |
| s3_bucket | s3_bucket | The S3 bucket |
| scale_out | boolean | Denotes a Flexgroup. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |
| storage_service | storage_service | Determines the placement of the volume that is to be provisioned. |
| tiering | tiering | |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning the volume. Only mirrored aggregates are used if this parameter is set to *true* and only unmirrored aggregates are used if this parameter is set to *false*. The default value is *true* for a MetroCluster configuration and is *false* for a non-MetroCluster configuration. |

container

| Name | Type | Description |
|---|---|---|
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| svm | svm | The SVM in which the container is located. |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning the volume. Only mirrored aggregates are used if this parameter is set to *true* and only unmirrored aggregates are used if this parameter is set to *false*. The default value is *true* for a MetroCluster configuration and is *false* for a non-MetroCluster configuration. |
| volumes | array[volumes] | A list of NAS volumes to provision. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Retrieve application templates

GET `/application/templates`

**Introduced In:** 9.6

Retrieves application templates.

### Query examples

The most useful queries on this API allows searches by name or protocol access. The following query returns all templates that are used to provision an Oracle application.

```
GET /application/templates?name=ora*
```

Similarly, the following query returns all templates that support SAN access.

```
GET /application/templates?protocol=san
```

### Learn more

- DOC /application

### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | query | False | Filter by name |
| protocol | string | query | False | Filter by protocol |
| description | string | query | False | Filter by description |
| missing_prerequisites | string | query | False | Filter by missing_prerequisites |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[application_template] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "string",
      "missing_prerequisites": "string",
      "name": "string",
      "nas": {
        "application_components": [
          {}
        ],
        "cifs_access": [
          {
            "access": "string",
            "user_or_group": "string"
          }
        ],
        "cifs_share_name": "string",
        "exclude_aggregates": [
          {
            "name": "string",
            "uuid": "string"
          }
        ],
        "nfs_access": [
          {
            "access": "string",
            "host": "string"
          }
        ],
        "protection_type": {
```

```
          "local_policy": "string",
          "local_rpo": "string",
          "remote_rpo": "string"
        }
      },
      "nvme": {
        "components": [
          {
            "name": "string",
            "os_type": "string",
            "performance": {
              "storage_service": {
                "name": "string"
              }
            },
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            },
            "subsystem": {
              "hosts": [
                {
                  "dh_hmac_chap": {
                    "group_size": "string",
                    "hash_function": "string"
                  },
                  "nqn": "string",
                  "priority": "string"
                }
              ],
              "name": "string",
              "os_type": "string",
              "uuid": "string"
            },
            "tiering": {
              "policy": "string"
            }
          }
        ],
        "os_type": "string",
        "rpo": {
          "local": {
            "name": "string",
            "policy": "string"
```

```
          },
          "remote": {
            "name": "string"
          }
        }
      },
      "protocol": "string",
      "s3_bucket": {
        "application_components": [
          {
            "access_policies": [
              {
                "actions": [
                  "string"
                ],
                "conditions": [
                  {
                    "delimiters": [
                      "string"
                    ],
                    "max_keys": [
                      "integer"
                    ],
                    "operator": "string",
                    "prefixes": [
                      "string"
                    ],
                    "source_ips": [
                      "string"
                    ],
                    "usernames": [
                      "string"
                    ]
                  }
                ],
                "effect": "string",
                "principals": [
                  "string"
                ],
                "resources": [
                  "string"
                ],
                "sid": "string"
              }
            ],
            "bucket_endpoint_type": "string",
```

```
            "comment": "string",
            "default_retention_period": "string",
            "exclude_aggregates": [
              {
                "name": "string",
                "uuid": "string"
              }
            ],
            "name": "string",
            "nas_path": "string",
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            },
            "retention_mode": "string",
            "snapshot_policy": {
              "name": "string",
              "uuid": "string"
            },
            "storage_service": {
              "name": "string"
            },
            "uuid": "string",
            "versioning_state": "string"
          }
        ],
        "protection_type": {
          "remote_rpo": "string"
        }
      },
      "san": {
        "application_components": [
          {
            "igroup_name": "string",
            "name": "string",
            "os_type": "string",
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            },
            "storage_service": {
              "name": "string"
```

```
            },
            "tiering": {
              "policy": "string"
            }
          }
        ],
        "exclude_aggregates": [
          {
            "name": "string",
            "uuid": "string"
          }
        ],
        "os_type": "string",
        "protection_type": {
          "local_policy": "string",
          "local_rpo": "string",
          "remote_rpo": "string"
        }
      }
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| element_count | integer | The number of storage elements (LUNs for SAN) of the database to maintain. Must be an even number between 2 and 16. Odd numbers will be rounded up to the next even number within range. |
| replication_factor | integer | The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

initiator_objects

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

protection_type

| Name | Type | Description |
|---|---|---|
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

secondary_igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
|---|---|---|
| dataset | dataset | |
| os_type | string | The name of the host OS running the application. |
| primary_igroup_name | string | The initiator group for the primary. |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

| Name | Type | Description |
|---|---|---|
| id | integer | The ID of an existing NFS export policy. |

| Name | Type | Description |
|---|---|---|
| name | string | The name of an existing NFS export policy. |

component

| Name | Type | Description |
|---|---|---|
| name | string | Name of the source component. |

svm

| Name | Type | Description |
|---|---|---|
| name | string | Name of the source SVM. |

origin

| Name | Type | Description |
|---|---|---|
| component | component | |
| svm | svm | |

flexcache

| Name | Type | Description |
|---|---|---|
| dr_cache | boolean | Dr-cache is a FlexCache volume create time option that has the same flexgroup-msid as that of the origin of a FlexCache volume. By default, dr-cache is disabled. The flexgroup-msid of the FlexCache volume does not need to be same as that of the origin of a FlexCache volume. |
| origin | origin | |

policy

| Name | Type | Description |
|---|---|---|
| name | string | The name of an existing QoS policy. |
| uuid | string | The UUID of an existing QoS policy. Usage: <UUID> |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as quot;Plt;num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period. |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| retention | retention | |
| snaplock_type | string | The SnapLock type of the smart container. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

nas_application_components_tiering

application-components.tiering

| Name | Type | Description |
|---|---|---|
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
|---|---|---|
| export_policy | export_policy | |
| flexcache | flexcache | |
| name | string | The name of the application component. |
| qos | qos | |
| scale_out | boolean | Denotes a Flexgroup. |
| share_count | integer | The number of shares in the application component. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Indicates whether Snapshot copy locking is enabled on the volume. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

| Name | Type | Description |
|---|---|---|
| access | string | The CIFS access granted to the user or group. |
| user_or_group | string | The name of the CIFS user or group that will be granted access. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the aggregate to exclude. Usage: <aggregate name> |
| uuid | string | The ID of the aggregate to exclude. Usage: <UUID> |

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The NFS access granted. |
| host | string | The name of the NFS entity granted access. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

nas

A generic NAS application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| cifs_access | array[app_cifs_access] | The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access. |

| Name | Type | Description |
| --- | --- | --- |
| cifs_share_name | string | The name of the CIFS share. Usage: <Share> |
| exclude_aggregates | array[exclude_aggregates] | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

performance

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| qos | qos | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |

| Name | Type | Description |
|---|---|---|
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
|---|---|---|
| name | string | The local RPO of the application. |
| policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |

remote

| Name | Type | Description |
|---|---|---|
| name | string | The remote RPO of the application. |

rpo

| Name | Type | Description |
|---|---|---|
| local | local | |
| remote | remote | |

zapp_nvme

An NVME application.

| Name | Type | Description |
|---|---|---|
| components | array[components] | |
| os_type | string | The name of the host OS running the application. |
| rpo | rpo | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| mirrored | boolean | Specifies whether the redo log group should be mirrored. |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| ora_home | ora_home | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| copies | integer | The number of CRS volumes. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
|---|---|---|
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| operator | string | Policy Condition Operator. |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[zapp_s3_bucket_application_components_access_policies_conditions] | conditions. |
| effect | string | Allow or Deny Access. |
| principals | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| resources | array[string] | |
| sid | string | Statement Identifier Usage: <(size 1..256)> |

snapshot_policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the snapshot policy that is used for the S3 bucket. |
| uuid | string | The UUID of the snapshot policy that is used for the S3 bucket. Usage: <UUID> |

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
| --- | --- | --- |
| access_policies | array[zapp_s3_bucket_application_components_access_policies] | The list of S3 objectstore policies to be created. |
| bucket_endpoint_type | string | The type of bucket. |
| capacity_tier | boolean | Prefer lower latency storage under similar media costs. |
| comment | string | Object Store Server Bucket Description Usage: <(size 1..256)> |

| Name | Type | Description |
| --- | --- | --- |
| default_retention_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as quot;Plt;num>Y" and quot;Plt;num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported. Usage: {{<integer> days\|years} \| none} |
| exclude_aggregates | array[exclude_aggregates] | |
| name | string | The name of the application component. |
| nas_path | string | The path to which the bucket corresponds to. |
| qos | qos | |
| retention_mode | string | The lock mode of the bucket. <br>compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects. <br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. For s3 type buckets, the default value is no_lock. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| snapshot_policy | snapshot_policy | |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |
| versioning_state | string | Bucket Versioning State. For nas type buckets, this field is not set. For s3 type buckets, the default value is disabled. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| remote_rpo | string | The remote RPO of the application. |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |
| protection_type | protection_type | |

san_application_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| qos | qos | |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

san

A generic SAN application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| exclude_aggregates | array[exclude_aggregates] | |
| os_type | string | The name of the host OS running the application. |

| Name | Type | Description |
| --- | --- | --- |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

access

| Name | Type | Description |
| --- | --- | --- |
| installer | string | SQL installer admin user name. |
| service_account | string | SQL service account user name. |

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| access | access | |

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| size | integer | The size of the desktops. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

hyper_v_access

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | Hyper-V service account. |

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
|---|---|---|
| desktops | desktops | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
|---|---|---|
| count | integer | The number of datastores to support. |
| size | integer | The size of the datastore. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
|---|---|---|
| datastore | datastore | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
|---|---|---|
| datastore | datastore | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

application_template

Application templates

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| description | string | Description. |

| Name | Type | Description |
| --- | --- | --- |
| missing_prerequisites | string | Missing prerequisites. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Template name. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protocol | string | Access protocol. |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve an application template

GET `/application/templates/{name}`

**Introduced In:** 9.6

Retrieves an application template.

### Template properties

Each application template has a set of properties. These properties are always nested under a property with the same name as the template. For example, when using the `s3_bucket` template, the properties are found nested inside the `s3_bucket` property. The properties nested under the template property are all specific to the template. The model for the application template object includes all the available templates, but only the object that corresponds to the template's name is returned, and only one is provided in any application API.

The model of each template includes a description of each property and its allowed values or usage. Default values are also indicated when available. The template properties returned by this API include an example value for each property.

### Template prerequisites

Each template has a set of prerequisites required for its use. If any of these prerequisites are not met, the `missing_prerequisites` property indicates which prerequisite is missing.

### Learn more

- DOC /application

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Template Name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| description | string | Description. |
| missing_prerequisites | string | Missing prerequisites. |
| mongo_db_on_san | mongo_db_on_san | MongoDB using SAN. |
| name | string | Template name. |
| nas | nas | A generic NAS application. |
| nvme | zapp_nvme | An NVME application. |
| oracle_on_nfs | oracle_on_nfs | Oracle using NFS. |
| oracle_on_san | oracle_on_san | Oracle using SAN. |
| oracle_rac_on_nfs | oracle_rac_on_nfs | Oracle RAC using NFS. |
| oracle_rac_on_san | oracle_rac_on_san | Oracle RAC using SAN. |
| protocol | string | Access protocol. |
| s3_bucket | zapp_s3_bucket | A generic S3 bucket application. |
| san | san | A generic SAN application. |
| sql_on_san | sql_on_san | Microsoft SQL using SAN. |
| sql_on_smb | sql_on_smb | Microsoft SQL using SMB. |
| vdi_on_nas | vdi_on_nas | A VDI application using NAS. |

| Name | Type | Description |
| --- | --- | --- |
| vdi_on_san | vdi_on_san | A VDI application using SAN. |
| vsi_on_nas | vsi_on_nas | A VSI application using NAS. |
| vsi_on_san | vsi_on_san | A VSI application using SAN. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "string",
  "missing_prerequisites": "string",
  "name": "string",
  "nas": {
    "application_components": [
      {}
    ],
    "cifs_access": [
      {
        "access": "string",
        "user_or_group": "string"
      }
    ],
    "cifs_share_name": "string",
    "exclude_aggregates": [
      {
        "name": "string",
        "uuid": "string"
      }
    ],
    "nfs_access": [
      {
        "access": "string",
        "host": "string"
      }
    ],
    "protection_type": {
      "local_policy": "string",
      "local_rpo": "string",
      "remote_rpo": "string"
    }
  },
  "nvme": {
    "components": [
      {
        "name": "string",
        "os_type": "string",
        "performance": {
```

```
          "storage_service": {
            "name": "string"
          }
        },
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "subsystem": {
          "hosts": [
            {
              "dh_hmac_chap": {
                "group_size": "string",
                "hash_function": "string"
              },
              "nqn": "string",
              "priority": "string"
            }
          ],
          "name": "string",
          "os_type": "string",
          "uuid": "string"
        },
        "tiering": {
          "policy": "string"
        }
      }
    ],
    "os_type": "string",
    "rpo": {
      "local": {
        "name": "string",
        "policy": "string"
      },
      "remote": {
        "name": "string"
      }
    }
  },
  "protocol": "string",
  "s3_bucket": {
    "application_components": [
      {
        "access_policies": [
```

```
        {
          "actions": [
            "string"
          ],
          "conditions": [
            {
              "delimiters": [
                "string"
              ],
              "max_keys": [
                "integer"
              ],
              "operator": "string",
              "prefixes": [
                "string"
              ],
              "source_ips": [
                "string"
              ],
              "usernames": [
                "string"
              ]
            }
          ],
          "effect": "string",
          "principals": [
            "string"
          ],
          "resources": [
            "string"
          ],
          "sid": "string"
        }
      ],
      "bucket_endpoint_type": "string",
      "comment": "string",
      "default_retention_period": "string",
      "exclude_aggregates": [
        {
          "name": "string",
          "uuid": "string"
        }
      ],
      "name": "string",
      "nas_path": "string",
      "qos": {
```

```
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "retention_mode": "string",
        "snapshot_policy": {
          "name": "string",
          "uuid": "string"
        },
        "storage_service": {
          "name": "string"
        },
        "uuid": "string",
        "versioning_state": "string"
      }
    ],
    "protection_type": {
      "remote_rpo": "string"
    }
  },
  "san": {
    "application_components": [
      {
        "igroup_name": "string",
        "name": "string",
        "os_type": "string",
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "storage_service": {
          "name": "string"
        },
        "tiering": {
          "policy": "string"
        }
      }
    ],
    "exclude_aggregates": [
      {
        "name": "string",
        "uuid": "string"
      }
```

```
    ],
    "os_type": "string",
    "protection_type": {
      "local_policy": "string",
      "local_rpo": "string",
      "remote_rpo": "string"
    }
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the database. |

dataset

| Name | Type | Description |
|---|---|---|
| element_count | integer | The number of storage elements (LUNs for SAN) of the database to maintain. Must be an even number between 2 and 16. Odd numbers will be rounded up to the next even number within range. |
| replication_factor | integer | The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. |
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

igroups

initiator_objects

mongo_db_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

secondary_igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group for each secondary. |

mongo_db_on_san

MongoDB using SAN.

| Name | Type | Description |
| --- | --- | --- |
| dataset | dataset | |
| os_type | string | The name of the host OS running the application. |
| primary_igroup_name | string | The initiator group for the primary. |
| protection_type | protection_type | |
| secondary_igroups | array[secondary_igroups] | |

export_policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ID of an existing NFS export policy. |
| name | string | The name of an existing NFS export policy. |

component

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source component. |

svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the source SVM. |

origin

| Name | Type | Description |
| --- | --- | --- |
| component | component | |
| svm | svm | |

flexcache

| Name | Type | Description |
| --- | --- | --- |
| dr_cache | boolean | Dr-cache is a FlexCache volume create time option that has the same flexgroup-msid as that of the origin of a FlexCache volume. By default, dr-cache is disabled. The flexgroup-msid of the FlexCache volume does not need to be same as that of the origin of a FlexCache volume. |
| origin | origin | |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of an existing QoS policy. |
| uuid | string | The UUID of an existing QoS policy. Usage: <UUID> |

qos

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as quot;Plt;num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period. |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| retention | retention | |
| snaplock_type | string | The SnapLock type of the smart container. |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the application component. |

object_stores

nas_application_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | |
| flexcache | flexcache | |
| name | string | The name of the application component. |
| qos | qos | |
| scale_out | boolean | Denotes a Flexgroup. |
| share_count | integer | The number of shares in the application component. |
| snaplock | snaplock | |
| snapshot_locking_enabled | boolean | Indicates whether Snapshot copy locking is enabled on the volume. |
| storage_service | storage_service | |
| tiering | nas_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member shares. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

app_cifs_access

The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The CIFS access granted to the user or group. |
| user_or_group | string | The name of the CIFS user or group that will be granted access. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the aggregate to exclude. Usage: <aggregate name> |
| uuid | string | The ID of the aggregate to exclude. Usage: <UUID> |

app_nfs_access

The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.

| Name | Type | Description |
| --- | --- | --- |
| access | string | The NFS access granted. |
| host | string | The name of the NFS entity granted access. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| local_policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |
| local_rpo | string | The local RPO of the application. |
| remote_rpo | string | The remote RPO of the application. |

nas

A generic NAS application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| cifs_access | array[app_cifs_access] | The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access. |

| Name | Type | Description |
| --- | --- | --- |
| cifs_share_name | string | The name of the CIFS share. Usage: <Share> |
| exclude_aggregates | array[exclude_aggregates] | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

performance

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | |

dh_hmac_chap

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | Authentication Diffie-Hellman Group. |
| hash_function | string | Authentication Hash Function. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | |
| nqn | string | The host NQN. |
| priority | string | Host Priority. |

zapp_nvme_components_subsystem

components.subsystem

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subsystem accessing the component. If neither the name nor the UUID is provided, the name defaults to <application-name>_<component-name>, whether that subsystem already exists or not. |
| os_type | string | The name of the host OS accessing the component. The default value is the host OS that is running the application. |
| uuid | string | The UUID of an existing subsystem to be granted access to the component. |

zapp_nvme_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

components

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the application component. |
| namespace_count | integer | The number of namespaces supported per request, with a total limit of 1024 per volume. |
| os_type | string | The name of the host OS running the application. |
| performance | performance | |
| qos | qos | |
| subsystem | zapp_nvme_components_subsystem | components.subsystem |
| tiering | zapp_nvme_components_tiering | application-components.tiering |

| Name | Type | Description |
| --- | --- | --- |
| total_size | integer | The total size of the component, spread across member namespaces. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

local

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local RPO of the application. |
| policy | string | The Snapshot copy policy to apply to each volume in the smart container. This property is only supported for smart containers. Usage: <snapshot policy> |

remote

| Name | Type | Description |
| --- | --- | --- |
| name | string | The remote RPO of the application. |

rpo

| Name | Type | Description |
| --- | --- | --- |
| local | local | |
| remote | remote | |

zapp_nvme

An NVME application.

| Name | Type | Description |
| --- | --- | --- |
| components | array[components] | |
| os_type | string | The name of the host OS running the application. |
| rpo | rpo | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the archive log. |

archive_log

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the archive log. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

db

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the database. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the ORACLE_HOME storage volume. |

ora_home

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the redo log group. |

redo_log

| Name | Type | Description |
| --- | --- | --- |
| mirrored | boolean | Specifies whether the redo log group should be mirrored. |
| size | integer | The size of the redo log group. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_on_nfs

Oracle using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| protection_type | protection_type | |
| redo_log | redo_log | |

oracle_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_on_san

Oracle using SAN.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| ora_home | ora_home | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle grid binary storage volume. |

grid_binary

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle grid binary storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the Oracle CRS volume. |

oracle_crs

| Name | Type | Description |
| --- | --- | --- |
| copies | integer | The number of CRS volumes. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

oracle_rac_on_nfs

Oracle RAC using NFS.

| Name | Type | Description |
| --- | --- | --- |
| archive_log | archive_log | |
| db | db | |
| grid_binary | grid_binary | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| protection_type | protection_type | |
| redo_log | redo_log | |

db_sids

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

oracle_rac_on_san

Oracle RAC using SAN.

| Name | Type | Description |
|---|---|---|
| archive_log | archive_log | |
| db | db | |
| db_sids | array[db_sids] | |
| grid_binary | grid_binary | |
| ora_home | ora_home | |
| oracle_crs | oracle_crs | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| redo_log | redo_log | |

zapp_s3_bucket_application_components_access_policies_conditions

conditions

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | |
| max_keys | array[integer] | |
| operator | string | Policy Condition Operator. |
| prefixes | array[string] | |
| source_ips | array[string] | |
| usernames | array[string] | |

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[zapp_s3_bucket_application_components_access_policies_conditions] | conditions. |
| effect | string | Allow or Deny Access. |
| principals | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| resources | array[string] | |
| sid | string | Statement Identifier Usage: <(size 1..256)> |

snapshot_policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the snapshot policy that is used for the S3 bucket. |
| uuid | string | The UUID of the snapshot policy that is used for the S3 bucket. Usage: <UUID> |

zapp_s3_bucket_application_components

The list of application components to be created.

| Name | Type | Description |
| --- | --- | --- |
| access_policies | array[zapp_s3_bucket_application_components_access_policies] | The list of S3 objectstore policies to be created. |
| bucket_endpoint_type | string | The type of bucket. |
| capacity_tier | boolean | Prefer lower latency storage under similar media costs. |
| comment | string | Object Store Server Bucket Description Usage: <(size 1..256)> |

| Name | Type | Description |
| --- | --- | --- |
| default_retention_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as quot;Plt;num>Y" and quot;Plt;num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported. Usage: {{<integer> days\|years} \| none} |
| exclude_aggregates | array[exclude_aggregates] | |
| name | string | The name of the application component. |
| nas_path | string | The path to which the bucket corresponds to. |
| qos | qos | |
| retention_mode | string | The lock mode of the bucket. <br>compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects. <br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. For s3 type buckets, the default value is no_lock. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| snapshot_policy | snapshot_policy | |
| storage_service | storage_service | |
| uuid | string | Object Store Server Bucket UUID Usage: <UUID> |
| versioning_state | string | Bucket Versioning State. For nas type buckets, this field is not set. For s3 type buckets, the default value is disabled. |

protection_type

| Name | Type | Description |
| --- | --- | --- |
| remote_rpo | string | The remote RPO of the application. |

zapp_s3_bucket

A generic S3 bucket application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[zapp_s3_bucket_application_components] | The list of application components to be created. |
| protection_type | protection_type | |

san_application_components_tiering

application-components.tiering

| Name | Type | Description |
| --- | --- | --- |
| policy | string | The storage tiering type of the application component. |

application_components

| Name | Type | Description |
| --- | --- | --- |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| lun_count | integer | The number of LUNs in the application component. |
| name | string | The name of the application component. |
| os_type | string | The name of the host OS running the application. |
| qos | qos | |
| storage_service | storage_service | |
| tiering | san_application_components_tiering | application-components.tiering |
| total_size | integer | The total size of the application component, split across the member LUNs. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |

san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

san

A generic SAN application.

| Name | Type | Description |
| --- | --- | --- |
| application_components | array[application_components] | |
| exclude_aggregates | array[exclude_aggregates] | |
| os_type | string | The name of the host OS running the application. |

| Name | Type | Description |
| --- | --- | --- |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the DB. |

db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the log DB. |

log

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the log DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[string] | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the temp DB. |

temp_db

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The size of the temp DB. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

sql_on_san

Microsoft SQL using SAN.

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| log | log | |
| os_type | string | The name of the host OS running the application. |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

access

| Name | Type | Description |
| --- | --- | --- |
| installer | string | SQL installer admin user name. |
| service_account | string | SQL service account user name. |

sql_on_smb

Microsoft SQL using SMB.

| Name | Type | Description |
| --- | --- | --- |
| access | access | |

| Name | Type | Description |
| --- | --- | --- |
| db | db | |
| log | log | |
| protection_type | protection_type | |
| server_cores_count | integer | The number of server cores for the DB. |
| temp_db | temp_db | |

storage_service

| Name | Type | Description |
| --- | --- | --- |
| name | string | The storage service of the desktops. |

desktops

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of desktops to support. |
| size | integer | The size of the desktops. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

hyper_v_access

| Name | Type | Description |
| --- | --- | --- |
| service_account | string | Hyper-V service account. |

vdi_on_nas

A VDI application using NAS.

| Name | Type | Description |
| --- | --- | --- |
| desktops | desktops | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vdi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

vdi_on_san

A VDI application using SAN.

| Name | Type | Description |
|---|---|---|
| desktops | desktops | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

storage_service

| Name | Type | Description |
|---|---|---|
| name | string | The storage service of the datastore. |

datastore

| Name | Type | Description |
|---|---|---|
| count | integer | The number of datastores to support. |
| size | integer | The size of the datastore. Usage: {<integer>[KB\|MB\|GB\|TB\|PB]} |
| storage_service | storage_service | |

vsi_on_nas

A VSI application using NAS.

| Name | Type | Description |
|---|---|---|
| datastore | datastore | |
| hyper_v_access | hyper_v_access | |
| nfs_access | array[app_nfs_access] | The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access. |
| protection_type | protection_type | |

vsi_on_san_new_igroups

The list of initiator groups to create.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | |

vsi_on_san

A VSI application using SAN.

| Name | Type | Description |
|---|---|---|
| datastore | datastore | |
| hypervisor | string | The name of the hypervisor hosting the application. |
| igroup_name | string | The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup. |
| protection_type | protection_type | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Storage cluster metrics

### ONTAP REST API storage cluster metrics endpoints

#### Overview

Retrieves cluster-wide historical capacity metrics.

#### Examples

### Retrieving cluster-wide historical space metrics for the past day

```
# The API:
/api/storage/cluster/metrics

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/cluster/metrics?interval=1d&fields=**" -H "accept:
application/hal+json"

# The response:
Snipped most of the records
{
"records": [
  {
    "uuid": "03ca3ef6-8173-11ef-825d-005056bb2d85",
    "timestamp": "2024-10-07T03:17:45Z",
    "status": "ok",
    "duration": "PT15S",
    "available_space": 862216192,
    "used_space": 1810432,
    "size": 864026624,
    "_links": {
      "self": {
```

```
        "href": "/api/storage/cluster/metrics/2024-10-
07T03%3A17%3A45Z?fields=**"
      }
    }
  },
  {
    "uuid": "03ca3ef6-8173-11ef-825d-005056bb2d85",
    "timestamp": "2024-10-07T02:18:00Z",
    "status": "partial_no_data",
    "duration": "PT15S",
    "available_space": 0,
    "used_space": 0,
    "size": 0,
    "_links": {
      "self": {
        "href": "/api/storage/cluster/metrics/2024-10-
07T02%3A18%3A00Z?fields=**"
      }
    }
  }
],
"num_records": 240,
"_links": {
  "self": {
    "href": "/api/storage/cluster/metrics?fields=**"
  }
}
}
```


**Retrieving historical metrics within a time range defined by timestamp**

```
curl -X GET "https://<mgmt-
ip>/api/storage/cluster/metrics?fields=**&timestamp=2024-10-
07T03:36:45Z..2024-10-07T03:36:30Z" -H "accept: application/hal+json"

#### Response:
{
"records": [
{
  "uuid": "03ca3ef6-8173-11ef-825d-005056bb2d85",
  "timestamp": "2024-10-07T03:36:45Z",
  "status": "ok",
  "duration": "PT15S",
  "available_space": 862183424,
  "used_space": 1843200,
```

```
  "size": 864026624,
  "_links": {
    "self": {
      "href": "/api/storage/cluster/metrics/2024-10-
07T03%3A36%3A45Z?fields=**"
    }
  }
},
{
  "uuid": "03ca3ef6-8173-11ef-825d-005056bb2d85",
  "timestamp": "2024-10-07T03:36:30Z",
  "status": "ok",
  "duration": "PT15S",
  "available_space": 862183424,
  "used_space": 1843200,
  "size": 864026624,
  "_links": {
    "self": {
      "href": "/api/storage/cluster/metrics/2024-10-
07T03%3A36%3A30Z?fields=**"
    }
  }
}
],
"num_records": 2,
"_links": {
"self": {
  "href": "/api/storage/cluster/metrics?fields=**&timestamp=2024-10-
07T03:36:45Z..2024-10-07T03:36:30Z"
}
}
}
```


# Cloud

## ONTAP REST API Cloud endpoints

Manages cloud (object storage) targets

## Manage cloud targets

### Manage cloud targets

#### Retrieving a collection of cloud targets

The cloud targets GET API retrieves all cloud targets defined in the cluster.

#### Creating cloud targets

The cluster administrator tells ONTAP how to connect to a cloud target. The following pre-requisites must be met before creating an object store configuration in ONTAP.
A valid data bucket or container must be created with the object store provider. This assumes that the user has valid account credentials with the object store provider to access the data bucket.
The ONTAP node must be able to connect to the object store.

This includes:

- Fast, reliable connectivity to the object store.
- An inter-cluster LIF (logical interface) must be configured on the cluster. ONTAP verifies connectivity prior to saving this configuration information.
- If SSL/TLS authentication is required, then valid certificates must be installed.
- FabricPool license (required for all object stores except SGWS).

#### Deleting cloud targets

If a cloud target is used by an aggregate, then the aggregate must be deleted before the cloud target can be deleted.

### Retrieve cloud targets in the cluster

GET `/cloud/targets`

**Introduced In:** 9.6

Retrieves the collection of cloud targets in the cluster.

#### Related ONTAP commands

- `storage aggregate object-store config show`
- `snapmirror object-store config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |
| container | string | query | False | Filter by container |
| cap_url | string | query | False | Filter by cap_url |
| access_key | string | query | False | Filter by access_key |
| certificate_validation_enabled | boolean | query | False | Filter by certificate_validation_enabled |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| port | integer | query | False | Filter by port |
| owner | string | query | False | Filter by owner |
| snapmirror_use | string | query | False | Filter by snapmirror_use |
| uuid | string | query | False | Filter by uuid |
| use_http_proxy | boolean | query | False | Filter by use_http_proxy<br>• Introduced in: 9.7 |
| cluster.name | string | query | False | Filter by cluster.name<br>• Introduced in: 9.7 |
| cluster.uuid | string | query | False | Filter by cluster.uuid<br>• Introduced in: 9.7 |
| url_style | string | query | False | Filter by url_style<br>• Introduced in: 9.8 |
| server | string | query | False | Filter by server |
| used | integer | query | False | Filter by used |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| read_latency_warning_threshold | integer | query | False | Filter by read_latency_warning_threshold<br>• Introduced in: 9.15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| azure_account | string | query | False | Filter by azure_account |
| ssl_enabled | boolean | query | False | Filter by ssl_enabled |
| authentication_type | string | query | False | Filter by authentication_type |
| provider_type | string | query | False | Filter by provider_type |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| server_side_encryption | string | query | False | Filter by server_side_encryption<br>• Introduced in: 9.7 |
| scope | string | query | False | Filter by scope<br>• Introduced in: 9.12 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cloud_target] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_key": "string",
      "authentication_type": "string",
      "azure_account": "string",
      "cap_url":
"https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&missi
on=mymission&role=myrole",
      "cluster": {
        "name": "string",
        "uuid": "string"
      },
      "container": "bucket1",
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "string",
      "owner": "string",
      "provider_type": "string",
      "scope": "string",
      "server": "string",
      "server_side_encryption": "string",
      "snapmirror_use": "string",
```

```
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "url_style": "string",
      "used": 0,
      "uuid": "string"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| name | string | The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state. |
| uuid | string | The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state. |

ipspace

IPspace to use in order to reach the cloud target.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM

information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster".

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cloud_target

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |
| azure_account | string | Azure account |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |

| Name | Type | Description |
| --- | --- | --- |
| container | string | Data bucket/container name. For FabricLink, a wildcard character "*" can also be specified to indicate that all the buckets in an SVM can use the same target information. However, for containers other than ONTAP, an exact name should be specified.<br>• example: bucket1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| ipspace | ipspace | IPspace to use in order to reach the cloud target. |
| name | string | Cloud target name |
| owner | string | Owner of the target. Allowed values are FabricPool, SnapMirror or S3_SnapMirror. A target can be used by only one feature. |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoogleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. For FabricLink, AWS_S3, SGWS, S3_Compatible, S3EMU, LOOPBACK and ONTAP_S3 are allowed.<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Cloud target UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a cloud target**

POST `/cloud/targets`

**Introduced In:** 9.6

Creates a cloud target.

**Required properties**

- `name` - Name for the cloud target.
- `owner` - Owner of the target: *fabricpool*, *snapmirror*.
- `provider_type` - Type of cloud provider: *AWS_S3*, *Azure_Cloud*, *SGWS*, *IBM_COS*, *AliCloud*, *GoogleCloud*, *ONTAP_S3*.
- `server` - Fully qualified domain name of the object store server. Required when `provider_type` is one of the following: *SGWS*, *IBM_COS*, *AliCloud*.
- `container` - Data bucket/container name.
- `access_key` - Access key ID if `provider_type` is not *Azure_Cloud* and `authentication_type` is *key*.
- `secret_password` - Secret access key if `provider_type` is not *Azure_Cloud* and `authentication_type` is *key*.

- `azure_account` - Azure account if `provider_type` is *Azure_Cloud*.
- `azure_private_key` - Azure access key if `provider_type` is *Azure_Cloud*.
- `cap_url` - Full URL of the request to a CAP server for retrieving temporary credentials if `authentication_type` is *cap*.
- `snapmirror_use` - Use of the cloud target if `owner` is *snapmirror*: data, metadata.
- `azure_msi_token` - Azure Managed Service Identity (MSI) token if `owner` is *fabricpool* or *snapmirror*, `provider_type` is *Azure_Cloud*, `authentication_type` if specified must be *azure_msi* and platform is Azure Netapp Files.

**Recommended optional properties**

- `authentication_type` - Authentication used to access the target: *key*, *cap*, *ec2_iam*, *gcp_sa*, *azure_msi*.
- `ssl_enabled` - SSL/HTTPS enabled or disabled.
- `port` - Port number of the object store that ONTAP uses when establishing a connection.
- `ipspace` - IPspace to use in order to reach the cloud target.
- `use_http_proxy` - Use the HTTP proxy when connecting to the object store server.
- `azure_sas_token` - Shared access signature to grant limited access to Azure storage account resources.
- `svm.name` or `svm.uuid` - Name or UUID of SVM if `owner` is *snapmirror*.
- `read_latency_warning_threshold` - Latency threshold to determine when to issue a warning alert EMS for a GET request.

**Default property values**

- `authentication_type`
- *ec2_iam* - if running in Cloud Volumes ONTAP in AWS
- *gcp_sa* - if running in Cloud Volumes ONTAP in GCP
- *azure_msi* - if running in Cloud Volumes ONTAP in Azure or if running on Azure NetApp Files platform with a Managed Service Identity (MSI) token.
- *key* - in all other cases.
- `server`
- *s3.amazonaws.com* - if `provider_type` is *AWS_S3*
- *blob.core.windows.net* - if `provider_type` is *Azure_Cloud*
- *storage.googleapis.com* - if `provider_type` is *GoogleCloud*
- `ssl_enabled` - *true*
- `port`
- *443* if `ssl_enabled` is *true*
- *80* if `ssl_enabled` is *false* and `provider_type` is not *SGWS*
- *8084* if `ssl_enabled` is *false* and `provider_type` is *SGWS*

- `ipspace` - *Default*
- `certificate_validation_enabled` - *true*
- `ignore_warnings` - *false*
- `check_only` - *false*
- `use_http_proxy` - *false*
- `server_side_encryption`
- *none* - if `provider_type` is *ONTAP_S3*
- *sse_s3* - if `provider_type` is not *ONTAP_S3*
- `url_style`
- *path_style* - if `provider_type` is neither *AWS_S3* nor *AliCloud*
- *virtual_hosted_style* - if `provider_type` is either *AWS_S3 or _AliCloud_*

#### Related ONTAP commands

- `storage aggregate object-store config create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ignore_warnings | boolean | query | False | Specifies whether or not warning codes should be ignored. |
| check_only | boolean | query | False | Do not create the target configuration, only check that the POST request succeeds. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |

| Name | Type | Description |
| --- | --- | --- |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |
| azure_account | string | Azure account |
| azure_msi_token | string | Managed Service Identity (MSI) token required to authenticate to the Azure object store. This form of authentication is only supported on Azure NetApp Files. |
| azure_private_key | string | Azure access key |
| azure_sas_token | string | Shared access signature token to access Azure containers and blobs. |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |

| Name | Type | Description |
| --- | --- | --- |
| container | string | Data bucket/container name. For FabricLink, a wildcard character "*" can also be specified to indicate that all the buckets in an SVM can use the same target information. However, for containers other than ONTAP, an exact name should be specified.<br>• example: bucket1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| ipspace | ipspace | IPspace to use in order to reach the cloud target. |
| name | string | Cloud target name |
| owner | string | Owner of the target. Allowed values are FabricPool, SnapMirror or S3_SnapMirror. A target can be used by only one feature. |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |
| provider_type | string | Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoogleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. For FabricLink, AWS_S3, SGWS, S3_Compatible, S3EMU, LOOPBACK and ONTAP_S3 are allowed.<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| secret_password | string | Secret access key for AWS_S3 and other S3 compatible provider types. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Cloud target UUID |

**Example request**

```
{
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "azure_msi_token": "string",
  "azure_private_key": "string",
  "azure_sas_token": "string",
  "cap_url":
"https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&missi
on=mymission&role=myrole",
  "cluster": {
    "name": "string",
    "uuid": "string"
  },
  "container": "bucket1",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "string",
  "owner": "string",
  "provider_type": "string",
  "scope": "string",
  "secret_password": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
  "uuid": "string"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 786436 | VLDB is not running. |
| 786908 | Capability check failed. |
| 787015 | "encryption-key-id" is not expected when encrypt is false. |
| 787016 | An object store configuration with the same combination of server and container name already exists. |

| Error Code | Description |
| --- | --- |
| 787020 | Failed to check whether the FabricPool capability exists. |
| 787021 | The FabricPool capability is not enabled. |
| 787030 | Object store configuration creation is not supported in this cluster. |
| 787036 | Server name is invalid. A valid server name must be a fully qualified domain name. |
| 787037 | Cannot verify availability of the object store. |
| 787038 | Object store provider type requires a FabricPool license. |
| 787039 | Failed to retrieve FabricPool capacity license information. |
| 787044 | Cannot create object store configuration as the number of object stores in the cluster has reached maximum allowed limit. |
| 787054 | An object store configuration with the same combination of server, azure account and container name already exists. |
| 787064 | Object store server provider type does not match object store provider type. Use the provider type that matches the server. |
| 787065 | Certificate validation must be enabled for the object store provider. |
| 787066 | Certificate validation cannot be specified when the "-is-ssl-enabled" parameter is false. |
| 787068 | Disabling certificate validation requires an effective cluster version of ONTAP 9.4.0 or later. |
| 787071 | Object store configuration creation requires an effective cluster version of ONTAP 9.4.0 or later. |
| 787082 | Creating an object store configuration requires an effective cluster version of ONTAP 9.5.0 or later. |
| 787089 | The object store provider supports the virtual hosted-style, and the specified "-server" contains the container name. The container specified in the "-server" parameter must be the same as the name of the container specified in the "-container" parameter. |
| 787133 | Could not retrieve temporary credentials from the CAP server. |
| 787134 | Could not retrieve temporary credentials from the CAP server due to an unexpected response. |
| 787136 | Specifying "CAP" as the "-auth-type" requires an effective cluster version of ONTAP 9.5.0 or later. |

| Error Code | Description |
| --- | --- |
| 787148 | The clock on node is behind by more than the maximum of 5 minutes. |
| 787149 | The clock on node is ahead by more than the maximum of 5 minutes. |
| 787158 | An object store configuration with the same combination of server and container name already exists. |
| 787159 | An object store configuration with the same name already exists. |
| 787179 | One or more clusters in this MetroCluster configuration do not have an effective cluster version of ONTAP 9.7.0 or later. |
| 787184 | Using HTTP proxies with FabricPool requires an effective cluster version of ONTAP 9.7.0 or later. |
| 787185 | There is no HTTP proxy for IPspace. Refer to the "vserver http-proxy" man page for details. |
| 787188 | Object store configuration belongs to another cluster and cannot be modified from the local cluster, unless the cluster is in switchover mode. |
| 787189 | Object store configuration name must not have the "-mc" suffix when the configuration is created for a local cluster. To create an object store configuration which belongs to another cluster, the cluster must be in switchover mode and "-cluster" must be specified. |
| 787209 | Object store is not accessible from the partner cluster in a MetroCluster configuration. |
| 787216 | Cannot perform object store configuration operations on a cluster that is waiting for switchback. |
| 787222 | Object store connectivity check failed on partner cluster in MetroCluster configuration. Wait a few minutes, and try the operation again. |
| 787223 | Specifying "GCP_SA" as the "-auth-type" requires an effective cluster version of ONTAP 9.7.0 or later. |
| 787227 | Specifying "Azure_MSI" as the "-auth-type" requires an effective cluster version of ONTAP 9.7 or later. |
| 787228 | SSL is required for this configuration. |
| 787229 | Cannot perform operation as URL style is not supported with object store provider type. |
| 787233 | The hash key for enabling this FabricPool feature is not present on the cluster. |
| 787234 | The hash key provided for the node to enable this FabricPool feature is not valid. |

| Error Code | Description |
| --- | --- |
| 787254 | The parameter is not supported on this system. |
| 787257 | The parameter "-encryption-context" is only applicable for AWS object store provider. |
| 787262 | The "create-container" option is applicable only for an SGWS or ONTAP_S3 object store provider. |
| 787301 | ONTAP S3 Bucket is of type NAS and is not supported as a target for FabricPool. |
| 787302 | Cannot use HTTP port with "-is-ssl-enabled" set to true. |
| 787303 | Cannot use HTTPS port with "-is-ssl-enabled" set to false. |
| 787325 | FabricPool is not supported on this platform. |
| 787350 | Creating an object store configuration with a Managed Service Identity (MSI) token is only supported on Azure NetApp files. |
| 787351 | Internal Error. Invalid authentication type. |
| 787352 | Creating an object store configuration with a Managed Service Identity (MSI) token requires an effective cluster version of ONTAP 9.16.1 or later. |
| 787353 | Creating an object store configuration with a Managed Service Identity (MSI) token is not supported for this owner. |
| 787354 | The specified properties are mutually exclusive. |
| 787355 | *azure_msi* `authentication-type` is only supported with *Azure_Cloud* `provider-type`. |
| 45940761 | Hostname cannot be resolved. Check the spelling of the hostname and check the system DNS availability using the "vserver services name-service dns check" command. |
| 45940778 | Bucket already exists. |
| 139591795 | Object store configuration for S3 SnapMirror representing ONTAP S3 object store provider only supports \"*\" as the container name. |
| 139591796 | Object store configuration name for S3 SnapMirror representing ONTAP S3 object store provider must be of the format \"vserver:<destination_vserver_uuid>:<source_vserver_uuid>\\\".</source_vserver_uuid></destination_vserver_uuid> |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

| Name | Type | Description |
|---|---|---|
| name | string | The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state. |
| uuid | string | The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state. |

ipspace

IPspace to use in order to reach the cloud target.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster".

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cloud_target

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |
| azure_account | string | Azure account |
| azure_msi_token | string | Managed Service Identity (MSI) token required to authenticate to the Azure object store. This form of authentication is only supported on Azure NetApp Files. |
| azure_private_key | string | Azure access key |
| azure_sas_token | string | Shared access signature token to access Azure containers and blobs. |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |

| Name | Type | Description |
| --- | --- | --- |
| container | string | Data bucket/container name. For FabricLink, a wildcard character "*" can also be specified to indicate that all the buckets in an SVM can use the same target information. However, for containers other than ONTAP, an exact name should be specified.<br>• example: bucket1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| ipspace | ipspace | IPspace to use in order to reach the cloud target. |
| name | string | Cloud target name |
| owner | string | Owner of the target. Allowed values are FabricPool, SnapMirror or S3_SnapMirror. A target can be used by only one feature. |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoogleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. For FabricLink, AWS_S3, SGWS, S3_Compatible, S3EMU, LOOPBACK and ONTAP_S3 are allowed.<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| secret_password | string | Secret access key for AWS_S3 and other S3 compatible provider types. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Cloud target UUID |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a cloud target

DELETE `/cloud/targets/{uuid}`

**Introduced In:** 9.6

Deletes the cloud target specified by the UUID. This request starts a job and returns a link to that job.

#### Related ONTAP commands

- `storage aggregate object-store config delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cloud target UUID |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 787013 | Cannot delete object store configuration as the store is attached to one or more aggregates. |
| 787014 | Cannot delete object store configuration as there are objects in the store. |
| 787078 | Cannot delete object store configuration. |
| 787188 | Object store configuration belongs to another cluster and cannot be modified from the local cluster, unless the cluster is in switchover mode. |
| 787216 | Cannot perform object store configuration operations on a cluster that is waiting for switchback. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a cloud target

GET `/cloud/targets/{uuid}`

**Introduced In:** 9.6

Retrieves the cloud target specified by the UUID.

#### Related ONTAP commands

- `storage aggregate object-store config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cloud target UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |
| azure_account | string | Azure account |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |

| Name | Type | Description |
| --- | --- | --- |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |
| container | string | Data bucket/container name. For FabricLink, a wildcard character "*" can also be specified to indicate that all the buckets in an SVM can use the same target information. However, for containers other than ONTAP, an exact name should be specified.<br>• example: bucket1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| ipspace | ipspace | IPspace to use in order to reach the cloud target. |
| name | string | Cloud target name |
| owner | string | Owner of the target. Allowed values are FabricPool, SnapMirror or S3_SnapMirror. A target can be used by only one feature. |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoogleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. For FabricLink, AWS_S3, SGWS, S3_Compatible, S3EMU, LOOPBACK and ONTAP_S3 are allowed.<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Cloud target UUID |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "cap_url":
"https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&missi
on=mymission&role=myrole",
  "cluster": {
    "name": "string",
    "uuid": "string"
  },
  "container": "bucket1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "string",
  "owner": "string",
  "provider_type": "string",
  "scope": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
```

```
  "uuid": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| name | string | The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state. |
| uuid | string | The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state. |

ipspace

IPspace to use in order to reach the cloud target.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster".

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a cloud target

PATCH `/cloud/targets/{uuid}`

**Introduced In:** 9.6

Updates the cloud target specified by the UUID with the fields in the body. This request starts a job and returns a link to that job.

#### Related ONTAP commands

- `storage aggregate object-store config modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cloud target UUID |
| ignore_warnings | boolean | query | False | Specifies whether or not warnings should be ignored. |
| check_only | boolean | query | False | Do not modify the configuration, only check that the PATCH request succeeds. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |
| azure_account | string | Azure account |
| azure_msi_token | string | Managed Service Identity (MSI) token required to authenticate to the Azure object store. This form of authentication is only supported on Azure NetApp Files. |
| azure_private_key | string | Azure access key |
| azure_sas_token | string | Shared access signature token to access Azure containers and blobs. |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |
| name | string | Cloud target name |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| secret_password | string | Secret access key for AWS_S3 and other S3 compatible provider types. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Cloud target UUID |

**Example request**

```
{
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "azure_msi_token": "string",
  "azure_private_key": "string",
  "azure_sas_token": "string",
  "cap_url":
"https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&missi
on=mymission&role=myrole",
  "name": "string",
  "scope": "string",
  "secret_password": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
  "uuid": "string"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 786436 | VLDB is not running. |
| 786908 | Capability check failed. |
| 787016 | An object store configuration with the same combination of server and container name already exists. |
| 787036 | Server name is invalid. A valid server name must be a fully qualified domain name. |
| 787038 | Object store provider type requires a FabricPool license. |
| 787039 | Failed to retrieve FabricPool capacity license information. |
| 787054 | An object store configuration with the same combination of server, azure account and container name already exists. |
| 787064 | Object store server provider type does not match object store provider type. Use the provider type that matches the server. |
| 787065 | Certificate validation must be enabled for the object store provider. |
| 787066 | Certificate validation cannot be specified when the "-is-ssl-enabled" parameter is false. |

| Error Code | Description |
| --- | --- |
| 787068 | Disabling certificate validation requires an effective cluster version of ONTAP 9.4.0 or later. |
| 787089 | The object store provider supports the virtual hosted-style, and the specified "-server" contains the container name. The container specified in the "-server" parameter must be the same as the name of the container specified in the "-container" parameter. |
| 787133 | Could not retrieve temporary credentials from the CAP server. |
| 787134 | Could not retrieve temporary credentials from the CAP server due to an unexpected response. |
| 787136 | Specifying "CAP" as the "-auth-type" requires an effective cluster version of ONTAP 9.5.0 or later. |
| 787148 | The clock on the node is behind by more than the maximum of 5 minutes. |
| 787149 | The clock on the node is ahead by more than the maximum of 5 minutes. |
| 787184 | Using HTTP proxies with FabricPool requires an effective cluster version of ONTAP 9.7.0 or later. |
| 787185 | There is no HTTP proxy for IPspace. Refer to the "vserver http-proxy" man page for details. |
| 787188 | Object store configuration belongs to cluster and cannot be modified from the local cluster, unless the cluster is in switchover mode. |
| 787209 | Object store is not accessible from the partner cluster in a MetroCluster configuration. |
| 787216 | Cannot perform object store configuration operations on a cluster that is waiting for switchback. |
| 787223 | Specifying "GCP_SA" as the "-auth-type" requires an effective cluster version of ONTAP 9.7.0 or later. |
| 787227 | Specifying "Azure_MSI" as the "-auth-type" requires an effective cluster version of ONTAP 9.7 or later. |
| 787228 | SSL is required for this configuration. |
| 787229 | Cannot perform operation as URL style is not supported with object store provider type. |
| 787233 | The hash key for enabling this FabricPool feature is not present on the cluster. |
| 787234 | The hash key provided for the node to enable this FabricPool feature is not valid. |
| 787254 | The parameter is not supported on this system. |

| Error Code | Description |
| --- | --- |
| 787257 | The parameter "-encryption-context" is only applicable for AWS object store provider. |
| 787301 | ONTAP S3 Bucket is of type NAS and is not supported as a target for FabricPool. |
| 787302 | Cannot use HTTP port with "-is-ssl-enabled" set to true. |
| 787303 | Cannot use HTTPS port with "-is-ssl-enabled" set to false. |
| 787306 | Object store is not accessible from the partner cluster in a MetroCluster configuration. |
| 787325 | FabricPool is not supported on this platform. |
| 787350 | Modifying an object store configuration with a Managed Service Identity (MSI) token is only supported on Azure NetApp Files. |
| 787351 | Internal Error. Invalid authentication type. |
| 787352 | Modifying an object store configuration with a Managed Service Identity (MSI) token requires an effective cluster version of ONTAP 9.16.1 or later. |
| 787353 | Modifying an object store configuration with a Managed Service Identity (MSI) token is not supported for this owner. |
| 787354 | The specified properties are mutually exclusive. |
| 787355 | *azure_msi* `authentication-type` is only supported with *Azure_Cloud* `provider-type`. |
| 787394 | Modifying the "-new-object-store-name\" or "-server\" without providing a Managed Identity (MI) token is not supported for this configuration. |
| 45940761 | Hostname cannot be resolved. Check the spelling of the hostname and check the system DNS availability using the "vserver services name-service dns check" command. |
| 45940778 | Bucket already exists. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

ipspace

IPspace to use in order to reach the cloud target.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster".

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cloud_target

| Name | Type | Description |
|---|---|---|
| access_key | string | Access key ID for AWS_S3 and other S3 compatible provider types. |
| authentication_type | string | Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| azure_account | string | Azure account |
| azure_msi_token | string | Managed Service Identity (MSI) token required to authenticate to the Azure object store. This form of authentication is only supported on Azure NetApp Files. |
| azure_private_key | string | Azure access key |
| azure_sas_token | string | Shared access signature token to access Azure containers and blobs. |
| cap_url | string | This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store. |
| certificate_validation_enabled | boolean | Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types.<br>• Introduced in: 9.6<br>• x-nullable: true |
| cluster | cluster | |
| name | string | Cloud target name |
| port | integer | Port number of the object store that ONTAP uses when establishing a connection. Required in POST.<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| read_latency_warning_threshold | integer | The warning threshold for read latency that is used to determine when an alert ems for a read operation from an object store should be issued. |
| scope | string | If the cloud target is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| secret_password | string | Secret access key for AWS_S3 and other S3 compatible provider types. |
| server | string | Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region> |

| Name | Type | Description |
| --- | --- | --- |
| server_side_encryption | string | Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.<br>• enum: ["none", "sse_s3", "sse_kms", "dsse_kms"]<br>• Introduced in: 9.7<br>• x-nullable: true |
| snapmirror_use | string | Use of the cloud target by SnapMirror. |
| ssl_enabled | boolean | SSL/HTTPS enabled or not |
| svm | svm | This field is only applicable when used for SnapMirror and FabricLink. For POST and PATCH, SVM information is required for SnapMirror and FabricLink targets when the scope is svm and is not allowed for FabricPool targets. For GET, this field is not set if the scope of the cloud target is "cluster". |
| url_style | string | URL style used to access S3 bucket. |
| use_http_proxy | boolean | Use HTTP proxy when connecting to the object store. |
| used | integer | The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Cloud target UUID |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Cluster

## ONTAP REST API Cluster endpoints

These APIs enable you to perform a number of independent workflows, including:

- Creating the cluster
- Adding nodes to the cluster
- Managing cluster configuration data (including name, version, NTP servers, name servers, and DNS domains)
- Managing node configuration data (including node names, models, serial numbers, and HA group information)

- Discovering the nodes on the cluster network that can be added to the cluster
- Viewing and updating current and recent jobs
- Updating the cluster software

#### Pre-Cluster APIs

A few of the cluster APIs (namely, POST/OPTIONS on /api/cluster, GET/HEAD/OPTIONS on /api/cluster/nodes, and calls on /api/cluster/jobs) are allowed before the cluster is created. These APIs support creation of the cluster and monitoring of its progress. Any other cluster API used before the cluster is created will fail.

## Manage clusters

### Manage clusters

### Overview

You can use this API to create a cluster, update cluster-wide configurations, and retrieve the current configuration details.

#### Creating a cluster

You can create a new cluster by issuing a POST request to /cluster. Parameters are provided in the body of the POST request to configure cluster-wide settings and add nodes during the cluster setup.

### Fields used for creating a cluster

The fields used for the cluster APIs fall into the following categories:

- Required cluster-wide configuration
- Optional cluster-wide configuration

### Required cluster-wide configuration

The following fields are always required for any POST /cluster request:

- name
- password

### Optional cluster-wide configuration

The following fields are used to set up additional cluster-wide configurations:

- location
- contact
- dns_domains
- name_servers
- ntp_servers
- timezone

- license
- configuration_backup
- management_interface
- nodes
- active_directory

### Nodes field

The nodes field specifies the nodes to join to the cluster. To use this API, all nodes must run the same version of ONTAP. If you do not specify a node, the cluster is configured with one node added. The REST request is issued to the node that is added to the cluster. If you specify one node, do not use the "node.cluster_interface.ip.address" field. If you specify multiple nodes, specify the node to which the REST request is issued in addition to the remote nodes. Use the "node.cluster_interface.ip.address" field to identify each node. All other node fields are optional in all cases. If you provide a field for one node, you need to provide the same field for all nodes.

### Node networking fields

The cluster management interface and each node management interface use the cluster management interface subnet mask and gateway. For advanced configurations in which the cluster and node management interfaces are on different subnets, use the /network/ip/interface APIs to configure network interfaces after setup is complete.
The management interfaces are used to communicate with the name servers and NTP servers. The address family of the name servers and NTP servers must match the management interfaces address family.

### Single node cluster field

When the "single_node_cluster" field is set to "true", the cluster is created in single node cluster mode. You can provide a node field for this node for node-specific configuration but do not use the "node.cluster_interface.ip.address" field. Storage failover is configured to non-HA mode, and ports used for cluster ports are moved to the default IPspace. This might cause the node to reboot during setup. While a node reboots, the RESTful interface might not be available. See "Connection failures during cluster create" for more information.

### Create recommended aggregates parameter

When the "create_recommended_aggregates" parameter is set to "true", aggregates based on an optimal layout recommended by the system are created on each of the nodes in the cluster. The default setting is "false".

---

#### Performance monitoring

Performance of the cluster can be monitored by the `metric.*` and `statistics.*` fields. These fields show the performance of the cluster in terms of IOPS, latency and throughput. The `metric.*` fields denote an average, whereas the `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

---

#### Analytics auto-enable properties

New SVMs will use the values set for the "auto_enable_analytics" and "auto_enable_activity_tracking" fields as the default for new volumes. The default setting is false.

### Setting auto_enable_analytics

```
# The API:
/api/cluster/

# The Call:
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d '{
"auto_enable_analytics" : "true"}'

# The response:
{
"job": {
  "uuid": "a079cfd8-50d0-11ed-9a7f-005056acd56e",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/a079cfd8-50d0-11ed-9a7f-005056acd56e"
    }
  }
}
}
```

### Setting auto_enable_activity_tracking

```
# The API:
/api/cluster/

# Modify  field:
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d '{
"auto_enable_activity_tracking" : "true"}'

# The response:
{
"job": {
  "uuid": "8a8b29e6-557a-11ed-92cc-005056ac76ec",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8a8b29e6-557a-11ed-92cc-005056ac76ec"
    }
  }
}
}
```

---

#### Monitoring cluster create status

### Errors before the job starts

Configuration in the POST /cluster request is validated before the cluster create job starts. If an invalid configuration is found, an HTTP error code in the 4xx range is returned. No cluster create job is started.

### Polling on the job

After a successful POST /cluster request is issued, an HTTP error code of 202 is returned along with a job UUID and link in the body of the response. The cluster create job continues asynchronously and is monitored with the job UUID using the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

### Errors during the job

If a failure occurs during the cluster create job, the job body provides details of the error along with error code fields. See the error table under "Responses" in the POST /cluster documentation for common error codes and descriptions.

### Rerunning POST /cluster

The POST /cluster request can be rerun if errors occur. When rerunning the request, use the same body and query parameters. You can change the value of any field in the original body or query, but you cannot change the provided fields. For example, an initial request might have a body section as follows:
+

```
body =
{
"name": "clusCreateRerun",
"password": "openSesame",
"nodes": [
  {
    "cluster_interface": {
      "ip": {
        "address": "1.1.1.1"
      }
    }
  },
  {
    "cluster_interface": {
      "ip": {
        "address": "2.2.2.2"
      }
    }
  }
]
}
```

A rerun request updates the body details to:
+

```
body =
{
"name": "clusCreateRerun",
"password": "openSesame",
"nodes": [
  {
    "cluster_interface": {
      "ip": {
        "address": "3.3.3.3"
      }
    }
  },
  {
    "cluster_interface": {
      "ip": {
        "address": "4.4.4.4"
      }
    }
  }
]
}
```

A rerun request with the following body details is invalid:
+

```
body =
{
"name": "clusCreateRerun",
"password": "openSesame",
"nodes": [
  {
    "cluster_interface": {
      "ip": {
        "address": "3.3.3.3"
      }
    }
  }
]
}
```

Note that the password might already be configured. If a password is already configured and then a new password is provided, the new request overwrites the existing password. If a password is already configured either by another interface or by a previous POST request to /cluster, authenticate any future REST requests with that password. If a POST request to /cluster with the default return_timeout of 0 returns an error, then the password was not changed.

### Connection failures during cluster create

A request to poll the job status might fail during a cluster create job in the following two cases. In these cases, programmatic use of the RESTful interface might be resilient to these connection failures.

1. When the "single_node_cluster" flag is set to "true", the node might reboot. During this time, the RESTful interface might refuse connections and return errors on a GET request, or connection timeouts might occur. Programmatic use of the RESTful interface during reboots must consider these effects while polling a cluster create job.
2. The "mgmt_auto" LIF is removed during the cluster create job. A POST /cluster request might be issued on the "mgmt_auto" LIF. However, requests to poll the job status might fail during cluster create when the "mgmt_auto" LIF is removed. The "mgmt_auto" LIF is only removed if a cluster management interface is provided as an argument to POST /cluster, and only after the cluster management interface is created. Programmatic use of the POST /cluster API on the "mgmt_auto" LIF should be configured to dynamically switch to polling the job on the cluster management LIF.

---

#### Active Directory account for the cluster

An Active Directory account for the cluster can be retrieved, created, modified, and deleted by using the `active_directory.*` fields. An account can be created either during cluster creation using a POST request or after cluster creation using a PATCH request.
An Active Directory account can be deleted by passing `null` as the Active Directory name during a PATCH request. The `username` and `password` fields are required to create, modify, and delete an Active Directory account.
Creating a new Active Directory account for the cluster:

```
# The API:
/api/cluster/

# Modify  field:
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d '{ "active_directory":
{"username": "administrator", "password": "password", "name": "adaccount",
"fqdn": "test.com", "force_account_overwrite" : true, "security":
{"advertised_kdc_encryptions": ["des"]}}}'

# The response:
{
"job": {
  "uuid": "8a8b29e6-557a-11ed-92cc-005056ac76ec",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8a8b29e6-557a-11ed-92cc-005056ac76ec"
    }
  }
}
}
```

Deleting the Active Directory account of the cluster:

```
# The API:
/api/cluster/

# Modify  field:
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d '{ "active_directory":
{"username": "administrator","password": "password","name": null}}'

# The response:
{
"job": {
  "uuid": "8a8b29e6-557a-11ed-92cc-005056ac76ec",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8a8b29e6-557a-11ed-92cc-005056ac76ec"
    }
  }
}
}
```

---

#### Modifying cluster configurations

The following fields can be used to modify a cluster-wide configuration:

- name
- location
- contact
- dns_domains
- name_servers
- timezone
- auto-enable-analytics
- auto-enable-activity-tracking
- active_directory

---

### Examples

#### Minimally configuring a 2-node setup

```
# Body
minimal_2_node_cluster.txt(body):
{
"name": "clusCreateExample1",
"password": "openSesame",
"nodes": [
  {
    "cluster_interface": {
      "ip": {
        "address": "1.1.1.1"
      }
    }
  },
  {
    "cluster_interface": {
      "ip": {
        "address": "2.2.2.2"
      }
    }
  }
]
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster" -d
"@minimal_2_node_cluster.txt"
```

---

**Setting up a single node with additional node configuration and auto aggregate creation**

```
# Body
single_node_additional_config.txt(body):
{
"name": "clusCreateExample2",
"password": "openSesame",
"nodes": [
  {
    "name": "singleNode",
    "location": "Sunnyvale"
  }
]
}

# Request
curl -X POST "https://<mgmt-
ip>/api/cluster?single_node_cluster=true&create_recommended_aggregates=tru
e" -d "@single_node_additional_config.txt"
```

### Modifying a cluster-wide configuration

```
# Body
modify_cluster_config.txt(body):
{
"contact": "it@company.com"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d
"@modify_cluster_config.txt"
```

#### Creating a cluster using the cluster "create" operation

This example shows how to create a cluster using the cluster APIs. Specifically, this example shows the creation of a two-node cluster and uses information from the nodes themselves combined with user supplied information to configure the cluster.

### Preparing for setup

Before the REST APIs can be issued to create the cluster, the cluster must be wired up and powered on. The network connections between the nodes for the cluster interconnect and the connections to the management network must be completed. After the nodes are powered on, the nodes automatically configure interfaces on the platform’s default cluster ports to allow the nodes to discover each other during setup and expansion workflows. You must configure a management interface on one node or use the mgmt_auto LIF, which is

assigned an IP address using DHCP, to start using the REST APIs. By making a console connection to a node, the cluster setup wizard guides you through the configuration of the initial node management interface to which the REST calls can be sent. Once this step is completed, exit the wizard by typing "exit". You can then issue REST API requests.

1. Wire and power on the nodes.
2. Make a console connection to one node to access the cluster setup wizard.
3. Enter node management interface information to enable REST API requests to be sent to the node.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
  http://support.netapp.com/autosupport/
  Type yes to confirm and continue {yes}: yes
  Enter the node management interface port [e0c]:
    Enter the node management interface IP address: 10.224.82.249
    Enter the node management interface netmask: 255.255.192.0
    Enter the node management interface default gateway: 10.224.64.1
    A node management interface on port e0c with IP address 10.224.82.249
has been created.
    Use your web browser to complete cluster setup by accessing
    https://10.224.82.249
    Otherwise, press Enter to complete cluster setup using the command
line
    interface: exit
    Exiting the cluster setup wizard. Any changes you made have been
saved.
    The cluster administrator's account (username "admin") password is set
to the system default.
    Warning: You have exited the cluster setup wizard before completing
all
    of the tasks. The cluster is not configured. You can complete cluster
setup by typing
    "cluster setup" in the command line interface.
```

---

#### Discovering the nodes

If you issue a GET /api/cluster/nodes request when the nodes are not in a cluster, the API returns a list of nodes that were discovered on the cluster interconnect. Information returned includes the node's serial number, model, software version, UUID, and cluster interface address. The number of nodes returned should

be the same as the number of nodes expected to be in the cluster. If too many nodes are discovered, remove the nodes that should not be part of the cluster. If not enough nodes are discovered, verify all the nodes are powered on, that the connections to the cluster interconnect are complete, and retry the command.
+

```
# The API:
/api/cluster/nodes

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=state,uptime" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "6dce4710-c860-11e9-b5bc-005056bb6135",
    "name": "cluster1",
    "uptime": 134555,
    "state": "up",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6dce4710-c860-11e9-b5bc-005056bb6135"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=state,uptime"
  }
}
}
```

---

### Creating the cluster

When the node information is available, including each node's cluster interface address, you can assemble the information for creating the cluster. Provide the cluster name and the password for the admin account. The rest of the information is optional and can be configured later using other APIs. Provide the cluster interface address for each node to be included in the cluster so that you can connect to it while adding it to the cluster. In addition to the cluster interface address, you can provide the optional node name, location, and management interface information. If you do not provide node names, nodes are named based on the cluster name. The nodes' management interface subnet mask and gateway values are omitted and must be the same as the cluster management interface's subnet mask and gateway.
+

```
# The API:
/api/cluster

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster" -H "accept:
application/hal+json" -H "accept: application/hal+json" -d
'{"name":"cluster1","location":"datacenter1","contact":"me","dns_domains":
["example.com"],"name_servers":["10.224.223.130","10.224.223.131","10.224.
223.132"],"ntp_servers":["time.nist.gov"],"active_directory": {"username":
"administrator","password": "password","name": "adaccount","fqdn":
"test.com","force_account_overwrite" : true},
"management_interface":{"ip":{"address":"10.224.82.25","netmask":"255.255.
192.0","gateway":"10.224.64.1"}},"password":"mypassword","license":{"keys"
:["AMEPOSOIKLKGEEEEDGNDEKSJDEEE"]},"nodes":[{"cluster_interface":{"ip":{"a
ddress":"169.254.245.113"}},"name":"node1","management_interface":{"ip":{"
address":"10.224.82.29"}}},{"cluster_interface":{"ip":{"address":"169.254.
217.95"}},"name":"node2","management_interface":{"ip":{"address":"10.224.8
2.31"}}}]}'

# The response:
{
"job": {
  "uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
    }
  }
}
}
```

#### Monitoring the progress of cluster creation

To monitor the progress of the cluster create operation, poll the returned job link until the state value is no longer "running" or "queued".
+

```
# The API:
/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5bc07e2-1e9-11e9-a751-
005056bbd95f" -H "accept: application/hal+json"

# The response:
{
"uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
"description": "POST /api/cluster",
"state": "success",
"message": "success",
"code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
  }
}
}
```

#### Verifying the cluster information

After the cluster is created, you can verify the information applied using a number of APIs. You can retrieve most of the information provided using the /api/cluster and /api/cluster/nodes APIs. In addition, you can view the network interface and route information using the /api/network APIs. The following example shows how to retrieve the cluster information:
+

```
# The API:
/api/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster" -H "accept:
application/hal+json"

# The response:
{
"name": "C1_sti44-vsim-ucs515w_1621957038",
"uuid": "5f7f57c7-bd67-11eb-95f4-005056a7b9b1",
"location": "sti",
"contact": "example_name",
"version": {
  "full": "NetApp Release 9.10.1: Mon May 24 08:07:35 UTC 2021",
```

```
  "generation": 9,
  "major": 10,
  "minor": 1
},
"dns_domains": [
  "example.com"
],
"name_servers": [
  "192.0.2.1",
  "192.0.2.2"
],
"ntp_servers": [
  "192.0.2.3"
],
"peering_policy": {
  "minimum_passphrase_length": 8,
  "authentication_required": true,
  "encryption_required": false
},
"management_interfaces": [
  {
    "uuid": "beef2db7-bd67-11eb-95f4-005056a7b9b1",
    "name": "clus_mgmt",
    "ip": {
      "address": "192.0.2.4"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/beef2db7-bd67-11eb-95f4-
005056a7b9b1"
      }
    }
  },
  {
    "uuid": "cb63e02c-bd72-11eb-95f4-005056a7b9b1",
    "name": "sti44-vsim-ucs515w_cluster_mgmt_inet6",
    "ip": {
      "address": "2001:db8:ef56:gh78::ij90"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/cb63e02c-bd72-11eb-95f4-
005056a7b9b1"
      }
    }
  },
```

```
  {
    "uuid": "ea13dec1-bd72-11eb-bd00-005056a7f50e",
    "name": "sti44-vsim-ucs515x_cluster_mgmt_inet6",
    "ip": {
      "address": "2001:db8:ef56:gh78::ij91"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/ea13dec1-bd72-11eb-bd00-
005056a7f50e"
      }
    }
  }
],
"metric": {
  "timestamp": "2021-05-26T20:36:15Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2021-05-26T20:36:25Z",
  "status": "ok",
  "latency_raw": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
```

```
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"timezone": {
  "name": "America/New_York"
},
"san_optimized": false,
"active_directory": {
  "name": "ADACCOUNT",
  "fqdn": "TEST.COM",
  "organizational_unit": "CN=Computers"
},
"_links": {
  "self": {
    "href": "/api/cluster"
  }
}
}

# The API:
/api/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster?fields=management_interfaces"
-H "accept: application/hal+json"

# The response:
{
"management_interfaces": [
  {
    "uuid": "c661725a-19e9-11e9-a751-005056bbd95f",
    "name": "cluster_mgmt",
    "ip": {
      "address": "192.0.2.1"
    },
    "_links": {
```

```
      "self": {
        "href": "/api/network/ip/interfaces/c661725a-19e9-11e9-a751-
005056bbd95f"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster"
  }
}
}
```


### Retrieve the cluster configuration

GET `/cluster`

**Introduced In:** 9.6

Retrieves the cluster configuration.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| active_directory | active_directory | |
| auto_enable_activity_tracking | boolean | Indicates how new SVMs will default "auto_enable_activity_tracking" for new volumes. |

| Name | Type | Description |
| --- | --- | --- |
| auto_enable_analytics | boolean | Indicates how new SVMs will default "auto_enable_analytics" for new volumes. |
| certificate | certificate | Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests. |
| cluster_network_overlay_enabled | boolean | Indicates whether the cluster network overlay is enabled. |
| contact | string | |
| disaggregated | boolean | Specifies whether the cluster is designed for disaggregated storage. |
| dns_domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |

| Name | Type | Description |
| --- | --- | --- |
| location | string | |
| management_interfaces | array[management_interfaces] | |
| metric | metric | Performance numbers, such as IOPS latency and throughput. |
| name | string | |
| name_servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| ntp_servers | array[string] | Host name, IPv4 address, or IPv6 address for the external NTP time servers. |
| peering_policy | peering_policy | |
| san_optimized | boolean | Specifies if this cluster is an All SAN Array. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| timezone | timezone | Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:<br>• console messages;<br>• logging to internal ONTAP log files; and<br>• localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.<br>• Introduced in: 9.7 |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "active_directory": {
    "force_account_overwrite": "",
    "fqdn": "server1.com",
    "name": "account1",
    "organizational_unit": "CN=Test",
    "security": {
      "advertised_kdc_encryptions": [
        "string"
      ]
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "location": "building 1",
  "management_interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
```

```
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "cluster1",
  "name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "ntp_servers": [
    "time.nist.gov",
    "10.98.19.20",
    "2610:20:6F15:15::27"
  ],
  "peering_policy": {
    "minimum_passphrase_length": 0
  },
  "statistics": {
```

```
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "timezone": {
    "name": "America/New_York"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9241607 | Only POST/OPTIONS on /api/cluster, GET/HEAD/OPTIONS on /api/cluster/nodes, or calls on /api/cluster/jobs are available in precluster. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

active_directory_security

| Name | Type | Description |
|---|---|---|
| advertised_kdc_encryptions | array[string] | |

active_directory

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account is created. |
| security | active_directory_security | |

certificate

Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

configuration_backup

| Name | Type | Description |
| --- | --- | --- |
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

license

License keys or NLF contents.

| Name | Type | Description |
| --- | --- | --- |
| keys | array[string] | |

ip

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

| Name | Type | Description |
|---|---|---|
| ip | ip | Object to setup an interface along with its default router. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

management_interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

| Name | Type | Description |
|---|---|---|
| timestamp | string | The timestamp of the performance data. |

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

cluster_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
|---|---|---|
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
|---|---|---|
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
|---|---|---|
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
|---|---|---|
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
| --- | --- | --- |
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
| --- | --- | --- |
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
| --- | --- | --- |
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |

| Name | Type | Description |
| --- | --- | --- |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
| --- | --- | --- |
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
| --- | --- | --- |
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
| --- | --- | --- |
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |
| local | local | |
| partner | partner | |

hw_assist

The hardware assist information.

| Name | Type | Description |
|---|---|---|
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
|---|---|---|
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |

| Name | Type | Description |
| --- | --- | --- |
| autoupdate_enabled | boolean | Indicates whether the service processor can be automatically updated from ONTAP.<br>• Introduced in: 9.10<br>• x-ntap-readModify: true<br>• x-nullable: true |
| backup | backup | Provides the properties of the service processor backup partition. |
| dhcp_enabled | boolean | Set to "true" to use DHCP to configure an IPv4 interface. Do not provide values for address, netmask and gateway when set to "true". |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| ipv6_interface | ipv6_interface | Object to setup an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |

| Name | Type | Description |
| --- | --- | --- |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |

| Name | Type | Description |
|---|---|---|
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
|---|---|---|
| provider_type | string | Cloud provider where the VM is hosted. |

nodes

Complete node information

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |

| Name | Type | Description |
| --- | --- | --- |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| hw_assist | hw_assist | The hardware assist information. |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metric | metric | CPU performance for the nodes. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| statistics | statistics | Raw CPU performance for the nodes. |

| Name | Type | Description |
| --- | --- | --- |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |
| vm | vm | |

peering_policy

| Name | Type | Description |
| --- | --- | --- |
| authentication_required | boolean | Indicates whether authentication is required in the communication between cluster peers. If true, authentication is required to establish communication between cluster peers. |
| encryption_required | boolean | Indicates whether encryption is required in the communication between cluster peers. If true, encryption is required to establish communication between cluster peers. |
| minimum_passphrase_length | integer | Minimum required length for a passphrase. For more information on password strength best practices, see: https://cheatsheetseries.owasp.org/cheatsheets/Authentication_Cheat_Sheet.html#implement-proper-password-strength-controls |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS

value to calculate the average latency per I/O operation.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each

node’s:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information.
  Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols
  such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.<br><br>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.<br><br>**IANA time zone format**<br><br>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.<br><br>"Location" represents a compound name using additional forward slashes.<br><br>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".<br><br>**ONTAP traditional time zone**<br><br>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.<br><br>• example: America/New_York<br>• Introduced in: 9.7<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the cluster configuration

PATCH `/cluster`

**Introduced In:** 9.6

Updates the cluster configuration after the cluster is created.

#### Related ONTAP commands

- `cluster identity modify`
- `system node modify`
- `vserver services dns modify`
- `vserver services name-service dns modify`
- `timezone`
- `vserver active-directory create`
- `vserver active-directory modify`
- `vserver active-directory delete`
- `vserver cifs security modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| active_directory | active_directory | |
| auto_enable_activity_tracking | boolean | Indicates how new SVMs will default "auto_enable_activity_tracking" for new volumes. |

| Name | Type | Description |
| --- | --- | --- |
| auto_enable_analytics | boolean | Indicates how new SVMs will default "auto_enable_analytics" for new volumes. |
| certificate | certificate | Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests. |
| contact | string | |
| disaggregated | boolean | Specifies whether the cluster is designed for disaggregated storage. |
| dns_domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| location | string | |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| name_servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| peering_policy | peering_policy | |
| san_optimized | boolean | Specifies if this cluster is an All SAN Array. |
| timezone | timezone | Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:<br>• console messages;<br>• logging to internal ONTAP log files; and<br>• localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.<br>• Introduced in: 9.7 |
| uuid | string | |

**Example request**

```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "active_directory": {
    "force_account_overwrite": "",
    "fqdn": "server1.com",
    "name": "account1",
    "organizational_unit": "CN=Test",
    "password": "testpwd",
    "security": {
      "advertised_kdc_encryptions": [
        "string"
      ]
    },
    "username": "admin"
  },
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cluster_network_overlay_enabled": true,
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "location": "building 1",
  "management_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "cluster1",
  "name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "peering_policy": {
    "minimum_passphrase_length": 0
  },
```

```
  "timezone": {
    "name": "America/New_York"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655414 | Failed to the remove Active Directory machine. |
| 655431 | Username cannot be empty. |
| 655432 | Password cannot be empty. |
| 655562 | The NetBIOS name cannot be longer than 15 characters. |

| Error Code | Description |
| --- | --- |
| 655563 | NetBIOS name contains characters that are not allowed. |
| 655914 | Failed to create the Active Directory machine account. |
| 655915 | A CIFS server for this SVM already exists. Having both a CIFS server and an Active Directory account for the same SVM is not supported. Use the \"vserver cifs delete\" command to delete the existing CIFS server and try the command again. |
| 656464 | Failed to create the Active Directory machine account. Reason: Invalid Credentials. |
| 656465 | Failed to create the Active Directory machine account. Reason: An account with this name already exists. |
| 656466 | Failed to create the Active Directory machine account. Reason: Unable to connect to any domain controllers. |
| 656467 | Failed to create the Active Directory machine account. Reason: Organizational-Unit not found. |
| 656478 | Failed to create the Active Directory machine account. Reason: KDC has no support for encryption type. |
| 656483 | Active Directory account creation for the admin SVM requires an effective cluster version of 9.16.0 or later. |
| 3604491 | Updating timezone failed. |
| 3604520 | Internal error. System state is not correct to read or change timezone. |
| 8847361 | Too many DNS domains provided. |
| 8847362 | Too many name servers provided. |
| 8847400 | The "dns_domains" field is required when "name_servers" is specified. |
| 9240587 | A name must be provided. |
| 9240588 | The name is too long. |
| 12451843 | Certificate does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

active_directory_security

| Name | Type | Description |
|---|---|---|
| advertised_kdc_encryptions | array[string] | |

active_directory

| Name | Type | Description |
|---|---|---|
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it is overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account is created. |
| password | string | Administrator password required for Active Directory account creation, modification, and deletion. |
| security | active_directory_security | |
| username | string | Administrator username required for Active Directory account creation, modification, and deletion. |

certificate

Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

configuration_backup

| Name | Type | Description |
| --- | --- | --- |
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

license

License keys or NLF contents.

| Name | Type | Description |
| --- | --- | --- |
| keys | array[string] | |

ip

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | Object to setup an interface along with its default router. |

ip

IP information

management_interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

cluster_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |

| Name | Type | Description |
| --- | --- | --- |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
|---|---|---|
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
|---|---|---|
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |

| Name | Type | Description |
| --- | --- | --- |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
|---|---|---|
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
|---|---|---|
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
|---|---|---|
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
|---|---|---|
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
|---|---|---|
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
|---|---|---|
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
|---|---|---|
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |

hw_assist

The hardware assist information.

| Name | Type | Description |
| --- | --- | --- |
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
| --- | --- | --- |
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |
| autoupdate_enabled | boolean | Indicates whether the service processor can be automatically updated from ONTAP.<br>• Introduced in: 9.10<br>• x-ntap-readModify: true<br>• x-nullable: true |
| backup | backup | Provides the properties of the service processor backup partition. |
| dhcp_enabled | boolean | Set to "true" to use DHCP to configure an IPv4 interface. Do not provide values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| ipv6_interface | ipv6_interface | Object to setup an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Cloud provider where the VM is hosted. |

nodes

Complete node information

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |

| Name | Type | Description |
| --- | --- | --- |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |

| Name | Type | Description |
|---|---|---|
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

peering_policy

| Name | Type | Description |
|---|---|---|
| authentication_required | boolean | Indicates whether authentication is required in the communication between cluster peers. If true, authentication is required to establish communication between cluster peers. |
| encryption_required | boolean | Indicates whether encryption is required in the communication between cluster peers. If true, encryption is required to establish communication between cluster peers. |
| minimum_passphrase_length | integer | Minimum required length for a passphrase. For more information on password strength best practices, see: https://cheatsheetseries.owasp.org/cheatsheets/Authentication_Cheat_Sheet.html#implement-proper-password-strength-controls |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information.
  Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols
  such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.<br><br>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.<br><br>**IANA time zone format**<br><br>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.<br><br>"Location" represents a compound name using additional forward slashes.<br><br>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".<br><br>**ONTAP traditional time zone**<br><br>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.<br><br>• example: America/New_York<br>• Introduced in: 9.7<br>• x-nullable: true |

cluster

Complete cluster information

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| active_directory | active_directory | |
| auto_enable_activity_tracking | boolean | Indicates how new SVMs will default "auto_enable_activity_tracking" for new volumes. |
| auto_enable_analytics | boolean | Indicates how new SVMs will default "auto_enable_analytics" for new volumes. |
| certificate | certificate | Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests. |
| contact | string | |
| disaggregated | boolean | Specifies whether the cluster is designed for disaggregated storage. |

| Name | Type | Description |
| --- | --- | --- |
| dns_domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| location | string | |
| management_interfaces | array[management_interfaces] | |
| name | string | |
| name_servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |
| peering_policy | peering_policy | |
| san_optimized | boolean | Specifies if this cluster is an All SAN Array. |

| Name | Type | Description |
| --- | --- | --- |
| timezone | timezone | Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:<br>• console messages;<br>• logging to internal ONTAP log files; and<br>• localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.<br>• Introduced in: 9.7 |
| uuid | string | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a cluster

POST `/cluster`

**Introduced In:** 9.6

Creates a cluster.

#### Required properties

- `name`
- `password`

#### Recommended optional properties

- `location`
- `contact`
- `dns_domains`
- `name_servers`
- `ntp_servers`
- `license`
- `configuration_backup`
- `management_interface`
- `nodes`
- `timezone`

#### Learn more

- DOC /cluster

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| single_node_cluster | boolean | query | False | Configures a single node cluster. All cluster ports are reassigned to the default network. The storage failover settings are configured to non-HA. The node reboots during this operation. |
| create_recommended_aggregates | boolean | query | False | Create aggregates based on an optimal layout recommended by the system.<br>• Introduced in: 9.7<br>• Default value: |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| keep_precluster_config | boolean | query | False | This is used to keep temporary configuration settings that allow initial setup including a node scoped certificate and possibly an automatically created node management interface. This is useful when creating a GUI that does not replace the node management interface using POST on /api/cluster, but instead creates the interface at another time. The certificate also relates to creating a web based GUI so that the certificate lasts through the entire workflow.<br>To remove the temporary configuration settings when a custom setup workflow is complete, set the remove_precluster_config query parameter in a PATCH on /api/cluster.<br>• Introduced in: 9.12<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| active_directory | active_directory | |
| auto_enable_activity_tracking | boolean | Indicates how new SVMs will default "auto_enable_activity_tracking" for new volumes. |

| Name | Type | Description |
| --- | --- | --- |
| auto_enable_analytics | boolean | Indicates how new SVMs will default "auto_enable_analytics" for new volumes. |
| cluster_network_overlay_enabled | boolean | Indicates whether the cluster network overlay is enabled. |
| configuration_backup | configuration_backup | |
| contact | string | |
| disaggregated | boolean | Specifies whether the cluster is designed for disaggregated storage. |
| dns_domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| license | license | License keys or NLF contents. |
| location | string | |

| Name | Type | Description |
| --- | --- | --- |
| management_interface | management_interface | The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration. |
| management_interfaces | array[management_interfaces] | |
| name | string | |
| name_servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| nodes | array[nodes] | |
| ntp_servers | array[string] | Host name, IPv4 address, or IPv6 address for the external NTP time servers. |
| password | string | Initial admin password used to create the cluster. |
| peering_policy | peering_policy | |
| san_optimized | boolean | Specifies if this cluster is an All SAN Array. |
| timezone | timezone | Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:<br>• console messages;<br>• logging to internal ONTAP log files; and<br>• localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.<br>• Introduced in: 9.7 |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

**Example request**

```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "active_directory": {
    "force_account_overwrite": "",
    "fqdn": "server1.com",
    "name": "account1",
    "organizational_unit": "CN=Test",
    "password": "testpwd",
    "security": {
      "advertised_kdc_encryptions": [
        "string"
      ]
    },
    "username": "admin"
  },
  "configuration_backup": {
    "password": "yourpassword",
    "rest_method": "post",
    "url": "http://10.224.65.198/backups",
    "username": "me"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "license": {
    "keys": [
      "<LICENSE-KEY>"
    ]
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "management_interfaces": [
```

```
      {
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "name": "cluster1",
    "name_servers": [
      "10.224.65.20",
      "2001:db08:a0b:12f0::1"
    ],
    "nodes": [
      {
        "anti_ransomware_version": "1.0",
        "cluster_interface": {
          "ip": {
            "address": "10.10.10.7"
          }
        },
        "cluster_interfaces": [
          {
            "name": "lif1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "controller": {
          "board": "System Board XXVIII",
          "cpu": {
            "count": 20,
            "firmware_release": "string",
            "processor": "string"
          },
          "failed_fan": {
            "count": 1,
            "message": {
              "code": "111411207",
              "message": "There are no failed fans."
            }
          },
          "failed_power_supply": {
            "count": 1,
            "message": {
              "code": "111411208",
              "message": "There are no failed power supplies."
            }
          },
          "flash_cache": [
```

```
          {
            "capacity": 1024000000000,
            "device_id": 0,
            "firmware_file": "X9170_O000Z6300NVM",
            "firmware_version": "NA05",
            "hardware_revision": "A1",
            "model": "X1970A",
            "part_number": "119-00207",
            "serial_number": "A22P5061550000187",
            "slot": "6-1",
            "state": "string"
          }
        ],
        "frus": [
          {
            "id": "string",
            "state": "string",
            "type": "string"
          }
        ],
        "memory_size": 1024000000,
        "over_temperature": "string"
      },
      "date": "2019-04-17 11:49:26 -0400",
      "external_cache": {
        "is_enabled": 1,
        "is_hya_enabled": 1,
        "is_rewarm_enabled": 1
      },
      "external_cache_bypass": {
        "enabled": 1,
        "large_read_ops_allow_percent": 100,
        "reset": 1
      },
      "ha": {
        "giveback": {
          "failure": {
            "code": 852126,
            "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
          },
          "state": "failed",
          "status": [
            {
              "aggregate": {
                "name": "aggr1",
```

```
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              },
              "error": {
                "code": "852126",
                "message": "string"
              },
              "state": "string"
            }
          ]
        },
        "interconnect": {
          "adapter": "MVIA-RDMA",
          "state": "string"
        },
        "partners": [
          {
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "ports": [
          {
            "number": 0,
            "state": "active"
          }
        ],
        "takeover": {
          "failure": {
            "code": 852130,
            "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
          },
          "state": "failed"
        },
        "takeover_check": {
          "reasons": [
            "string"
          ]
        },
        "type": "string"
      },
      "location": "rack 2 row 5",
      "management_interface": {
        "ip": {
          "address": "10.10.10.7"
        }
```

```
      },
      "management_interfaces": [
        {
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "membership": "string",
      "metrocluster": {
        "ports": [
          {
            "name": "e1b"
          }
        ],
        "type": "string"
      },
      "model": "FAS3070",
      "name": "node-01",
      "nvlog": {
        "backing_type": "string",
        "swap_mode": "string"
      },
      "nvram": {
        "battery_state": "string",
        "id": 0
      },
      "owner": "Example Corp",
      "serial_number": "4048820-60-9",
      "service_processor": {
        "api_service": {
          "port": 0
        },
        "auto_config": {
          "ipv4_subnet": "ipv4_mgmt",
          "ipv6_subnet": "ipv6_mgmt"
        },
        "autoupdate_enabled": true,
        "backup": {
          "state": "string",
          "version": "11.6"
        },
        "dhcp_enabled": true,
        "firmware_version": "string",
        "ipv4_interface": {
          "address": "10.10.10.7",
          "gateway": "10.1.1.1",
```

```
          "netmask": "255.255.0.0",
          "setup_state": "string"
        },
        "last_update_state": "string",
        "link_status": "string",
        "mac_address": "string",
        "primary": {
          "state": "string",
          "version": "11.6"
        },
        "ssh_info": {
          "allowed_addresses": [
            "10.10.10.7/24"
          ]
        },
        "state": "string",
        "type": "string"
      },
      "snaplock": {
        "compliance_clock_time": "2018-06-04 19:00:00 +0000"
      },
      "state": "string",
      "storage_configuration": "string",
      "system_aggregate": {
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "system_id": 92027651,
      "system_machine_type": "7Y56-CTOWW1",
      "uptime": 300536,
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
      "vendor_serial_number": 791603000068,
      "vm": {
        "provider_type": "string"
      }
    }
  ],
  "ntp_servers": [
    "time.nist.gov",
    "10.98.19.20",
    "2610:20:6F15:15::27"
  ],
  "password": "mypassword",
  "peering_policy": {
    "minimum_passphrase_length": 0
  },
```

```
  "timezone": {
    "name": "America/New_York"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262245 | The value provided is invalid. |

| Error Code | Description |
| --- | --- |
| 655562 | The NetBIOS name cannot be longer than 15 characters. |
| 655915 | A CIFS server for this SVM already exists. Having both a CIFS server and an Active Directory account for the same SVM is not supported. Use the \"vserver cifs delete\" command to delete the existing CIFS server and try the command again. |
| 656464 | Failed to create the Active Directory machine account. Reason: Invalid Credentials. |
| 656465 | Failed to create the Active Directory machine account. Reason: An account with this name already exists. |
| 656466 | Failed to create the Active Directory machine account. Reason: Unable to connect to any domain controllers. |
| 656467 | Failed to create the Active Directory machine account. Reason: Organizational-Unit not found. |
| 656483 | Active Directory account creation for the admin SVM requires an effective cluster version of 9.16.0 or later. |
| 1179812 | This node cannot be configured as a single node cluster. |
| 1179813 | Fields set for one node must be set for all nodes. |
| 1179817 | The IP address, subnet mask, and gateway must all be provided for cluster management interface. |
| 1179818 | The IP address and gateway must be of the same family. |
| 1179821 | An IP address and subnet mask conflicts with an existing entry. |
| 1179823 | An invalid netmask was provided. |
| 1179824 | An invalid gateway was provided. |
| 1179825 | All management and cluster config IP addresses must belong to the same address family. |
| 1182809 | Cluster creation is already complete. |
| 1182810 | This platform does not support cluster network overlay. |
| 2097165 | An NTP server could not be reached. |
| 7077919 | The minimum length for the new password does not meet the policy. |
| 8847361 | Too many DNS domains provided. |
| 8847362 | Too many name servers provided. |
| 8847394 | An invalid DNS domain was provided. |

| Error Code | Description |
| --- | --- |
| 8978433 | An invalid license key was provided. |
| 9240587 | A name must be provided. |
| 9240594 | An invalid name was provided. |
| 39387137 | The URL provided is invalid. |
| 131727360 | A node could not be added to the cluster. This is a generic code, see response message for details. |
| 131727388 | Hostnames for NTP servers cannot be used without DNS configured. |
| 131727389 | URL and username are required for configuration backup. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

active_directory_security

| Name | Type | Description |
|---|---|---|
| advertised_kdc_encryptions | array[string] | |

active_directory

| Name | Type | Description |
|---|---|---|
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it is overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account is created. |
| password | string | Administrator password required for Active Directory account creation, modification, and deletion. |
| security | active_directory_security | |
| username | string | Administrator username required for Active Directory account creation, modification, and deletion. |

certificate

Support for this field will be removed in a future release. Please use /api/cluster/web for this field. Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

configuration_backup

| Name | Type | Description |
| --- | --- | --- |
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

license

License keys or NLF contents.

| Name | Type | Description |
| --- | --- | --- |
| keys | array[string] | |

ip

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | Object to setup an interface along with its default router. |

ip

IP information

management_interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

| Name | Type | Description |
| --- | --- | --- |
| ip | node_setup_ip | The IP configuration for cluster setup. |

cluster_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
|---|---|---|
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
|---|---|---|
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
|---|---|---|
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |

| Name | Type | Description |
| --- | --- | --- |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |

| Name | Type | Description |
| --- | --- | --- |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
|---|---|---|
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
|---|---|---|
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
|---|---|---|
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
| --- | --- | --- |
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
| --- | --- | --- |
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
| --- | --- | --- |
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |

hw_assist

The hardware assist information.

| Name | Type | Description |
| --- | --- | --- |
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

| Name | Type | Description |
| --- | --- | --- |
| ip | node_setup_ip | The IP configuration for cluster setup. |

management_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
| --- | --- | --- |
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |
| backup | backup | Provides the properties of the service processor backup partition. |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |

| Name | Type | Description |
| --- | --- | --- |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |

| Name | Type | Description |
| --- | --- | --- |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Cloud provider where the VM is hosted. |

nodes

Complete node information

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interface | cluster_interface | The cluster network IP address of the node to be added. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |

| Name | Type | Description |
| --- | --- | --- |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interface | management_interface | The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes. |
| management_interfaces | array[management_interfaces] | |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |

| Name | Type | Description |
| --- | --- | --- |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

peering_policy

| Name | Type | Description |
| --- | --- | --- |
| authentication_required | boolean | Indicates whether authentication is required in the communication between cluster peers. If true, authentication is required to establish communication between cluster peers. |

| Name | Type | Description |
| --- | --- | --- |
| encryption_required | boolean | Indicates whether encryption is required in the communication between cluster peers. If true, encryption is required to establish communication between cluster peers. |
| minimum_passphrase_length | integer | Minimum required length for a passphrase. For more information on password strength best practices, see: https://cheatsheetseries.owasp.org/cheatsheets/Authentication_Cheat_Sheet.html#implement-proper-password-strength-controls |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each

node’s:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information.
  Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols
  such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.<br><br>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.<br><br>**IANA time zone format**<br><br>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.<br><br>"Location" represents a compound name using additional forward slashes.<br><br>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".<br><br>**ONTAP traditional time zone**<br><br>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.<br><br>• example: America/New_York<br>• Introduced in: 9.7<br>• x-nullable: true |

cluster

Complete cluster information

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| active_directory | active_directory | |
| auto_enable_activity_tracking | boolean | Indicates how new SVMs will default "auto_enable_activity_tracking" for new volumes. |
| auto_enable_analytics | boolean | Indicates how new SVMs will default "auto_enable_analytics" for new volumes. |
| cluster_network_overlay_enabled | boolean | Indicates whether the cluster network overlay is enabled. |
| configuration_backup | configuration_backup | |
| contact | string | |
| disaggregated | boolean | Specifies whether the cluster is designed for disaggregated storage. |

| Name | Type | Description |
| --- | --- | --- |
| dns_domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| license | license | License keys or NLF contents. |
| location | string | |
| management_interface | management_interface | The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration. |
| management_interfaces | array[management_interfaces] | |
| name | string | |
| name_servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |
| nodes | array[nodes] | |

| Name | Type | Description |
| --- | --- | --- |
| ntp_servers | array[string] | Host name, IPv4 address, or IPv6 address for the external NTP time servers. |
| password | string | Initial admin password used to create the cluster. |
| peering_policy | peering_policy | |
| san_optimized | boolean | Specifies if this cluster is an All SAN Array. |
| timezone | timezone | Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:<br>• console messages;<br>• logging to internal ONTAP log files; and<br>• localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.<br>• Introduced in: 9.7 |
| uuid | string | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve cluster chassis

### Retrieve cluster chassis

#### Overview

You can use the chassis GET API to retrieve all of the chassis information in the cluster.

#### Examples

### Retrieving a list of chassis from the cluster

The following example shows the response with a list of chassis in the cluster:

```
# The API:
/api/cluster/chassis

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "id": "021352005981",
    "_links": {
      "self": {
        "href": "/api/cluster/chassis/021352005981"
      }
    }
  },
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/cluster/chassis"
  }
}
}
```

#### Retrieving a specific chassis from the cluster

The following example shows the response of the requested chassis. If there is no chassis with the requested ID, an error is returned.

```
# The API:
/api/cluster/chassis/{id}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis/021352005981" -H
"accept: application/hal+json"

# The response:
{
"id": "021352005981",
"state": "ok",
```

```
"nodes": [
  {
    "name": "node-1",
    "uuid": "6ede364b-c3d0-11e8-a86a-00a098567f31",
    "position": "top",
    "usbs": {
      "supported": true,
      "enabled": true,
      "ports": [
        {
          "connected": false
        }
      ]
    },
    "pcis": {
      "cards": [
        {
          "slot": "0",
          "device": "Gigabit Ethernet I210",
          "info": "\t  e0M MAC Address:    d0:39:ea:3f:06:2b (auto-1000t-
fd-up) \n\t  e0S MAC Address:    d0:39:ea:3f:06:2c (auto-1000t-fd-up) \n\t
Device Type:        1533\n\t  Firmware Version:   3.25-0.0 0x800005D1\n"
        },
        {
          "slot": "0",
          "device": "Intel Lewisburg series chipset SATA Controller",
          "info": "\t  Additional Info: 0 (0xaaf00000)   \n\t
SHM2S86Q120GLM22NP FW1146 114473MB 512B/sect (SPG190108HJ)  \n"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6ede364b-c3d0-11e8-a86a-00a098567f31"
      }
    }
  }
],
"frus": [
  {
    "id": "PSU2",
    "type": "psu",
    "state": "ok"
  },
  {
    "id": "PSU1",
```

```
      "type": "psu",
      "state": "ok"
    },
    {
      "id": "Fan2",
      "type": "fan",
      "state": "ok"
    },
    {
      "id": "Fan3",
      "type": "fan",
      "state": "ok"
    },
    {
      "id": "Fan1",
      "type": "fan",
      "state": "ok"
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/chassis/021352005981"
    }
  }
  }
```


### Retrieve a collection of chassis

GET `/cluster/chassis`

**Introduced In:** 9.6

Retrieves a collection of chassis.

#### Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

#### Learn more

- DOC /cluster/chassis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| frus.id | string | query | False | Filter by frus.id |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| frus.type | string | query | False | Filter by frus.type |
| frus.state | string | query | False | Filter by frus.state |
| nodes.position | string | query | False | Filter by nodes.position<br>• Introduced in: 9.8 |
| nodes.usbs.ports.connected | boolean | query | False | Filter by nodes.usbs.ports.connected<br>• Introduced in: 9.9 |
| nodes.usbs.enabled | boolean | query | False | Filter by nodes.usbs.enabled<br>• Introduced in: 9.9 |
| nodes.usbs.supported | boolean | query | False | Filter by nodes.usbs.supported<br>• Introduced in: 9.9 |
| nodes.name | string | query | False | Filter by nodes.name |
| nodes.pcis.cards.slot | string | query | False | Filter by nodes.pcis.cards.slot<br>• Introduced in: 9.9 |
| nodes.pcis.cards.device | string | query | False | Filter by nodes.pcis.cards.device<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| nodes.pcis.cards.info | string | query | False | Filter by nodes.pcis.cards.info<br>• Introduced in: 9.9 |
| nodes.uuid | string | query | False | Filter by nodes.uuid |
| id | string | query | False | Filter by id |
| shelves.uid | string | query | False | Filter by shelves.uid |
| state | string | query | False | Filter by state |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[chassis] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "frus": [
        {
          "id": "string",
          "state": "string",
          "type": "string"
        }
      ],
      "id": "021352005981",
      "nodes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "pcis": {
            "cards": [
              {
                "device": "Intel Lewisburg series chipset SATA
Controller",
                "info": "Additional Info: 0 (0xaaf00000)
SHM2S86Q120GLM22NP FW1146 114473MB 512B/sect (SPG190108GW)",
                "slot": "0"
              }
            ]
          },
          "position": "top",
          "usbs": {
            "ports": [
              {}
            ]
```

```
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "shelves": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "uid": "7777841915827391056"
          }
        ],
        "state": "string"
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cards

| Name | Type | Description |
|---|---|---|
| device | string | The description of the PCI card. |
| info | string | The info string from the device driver of the PCI card. |
| slot | string | The slot where the PCI card is placed. This can sometimes take the form of "6-1" to indicate slot and subslot. |

pcis

| Name | Type | Description |
|---|---|---|
| cards | array[cards] | |

ports

| Name | Type | Description |
| --- | --- | --- |
| connected | boolean | Indicates whether or not the USB port has a device connected to it. |

usbs

The status of the USB ports on the controller.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether or not the USB ports are enabled. |
| ports | array[ports] | |
| supported | boolean | Indicates whether or not USB ports are supported on the current platform. |

nodes

List of nodes in chassis.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| pcis | pcis | |
| position | string | The position of the node in the chassis, when viewed from the rear of the system. |
| usbs | usbs | The status of the USB ports on the controller. |
| uuid | string | |

shelf_reference

Shelf

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uid | string | |

chassis

| Name | Type | Description |
|---|---|---|
| frus | array[frus] | List of FRUs in the chassis. |
| id | string | |
| nodes | array[nodes] | List of nodes in the chassis. |
| shelves | array[shelf_reference] | List of shelves in chassis. |
| state | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a chassis

GET `/cluster/chassis/{id}`

**Introduced In:** 9.6

Retrieves a specific chassis.

#### Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

#### Learn more

- DOC /cluster/chassis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | string | path | True | Chassis ID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| frus | array[frus] | List of FRUs in the chassis. |
| id | string | |
| nodes | array[nodes] | List of nodes in the chassis. |
| shelves | array[shelf_reference] | List of shelves in chassis. |
| state | string | |

**Example response**


```
{
  "frus": [
    {
      "id": "string",
      "state": "string",
      "type": "string"
    }
  ],
  "id": "021352005981",
  "nodes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "pcis": {
        "cards": [
          {
            "device": "Intel Lewisburg series chipset SATA Controller",
            "info": "Additional Info: 0 (0xaaf00000)
SHM2S86Q120GLM22NP FW1146 114473MB 512B/sect (SPG190108GW)",
            "slot": "0"
          }
        ]
      },
      "position": "top",
      "usbs": {
        "ports": [
          {}
        ]
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "shelves": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uid": "7777841915827391056"
```

```
      }
    ],
    "state": "string"
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cards

| Name | Type | Description |
|---|---|---|
| device | string | The description of the PCI card. |
| info | string | The info string from the device driver of the PCI card. |
| slot | string | The slot where the PCI card is placed. This can sometimes take the form of "6-1" to indicate slot and subslot. |

pcis

| Name | Type | Description |
|---|---|---|
| cards | array[cards] | |

ports

| Name | Type | Description |
|---|---|---|
| connected | boolean | Indicates whether or not the USB port has a device connected to it. |

usbs

The status of the USB ports on the controller.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the USB ports are enabled. |
| ports | array[ports] | |
| supported | boolean | Indicates whether or not USB ports are supported on the current platform. |

nodes

List of nodes in chassis.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| pcis | pcis | |
| position | string | The position of the node in the chassis, when viewed from the rear of the system. |
| usbs | usbs | The status of the USB ports on the controller. |
| uuid | string | |

shelf_reference

Shelf

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve cluster counter tables

### Retrieve cluster counter tables

#### Overview

The Counter Manager subsystem allows both manual and automated processes to access statistical information about various aspects of the ONTAP system. The information is most often utilized to assess the current performance of the system.

The data architecture is broken down into four components:

- Tables
- Rows
- Counters / Properties
- Aggregation

### Tables

A table represents a collection of statistics that are grouped according to a common feature or function. An example counter manager table is for network adapters. This table would contain statistics related to the network adapter's performance such as the number of packets, rate of flow and error counters.

A table is described by its schema which includes a detailed description about the various statistics included, their format and their purpose.

The table catalog is a collection of all the statistical tables that the ONTAP REST interface supports, which can be queried to find information about a data point of interest.

### Rows

Each table is populated with a list of rows. Each row is identified by a unique key and represents a specific statistical entity within the system. For example, a system may contain multiple network adapters that are represented by several records in the network adapter table.

### Counter / Property

A counter is the basic 'numeric' statistical unit of the architecture.

A property is the basic 'string' statistical unit of the architecture.

Counter values can be organized as singular values or into multi-dimensional arrays. An array can be one or two dimensional; formatted as a list of label / value pairs. Additional detail can be found in the "counter" model definition.

A table schema definition consists of multiple counters and properties.

Counters are classified according to their type. The available type options are the following:

- average
- rate
- raw
- delta
- percent

Average and percent counters specify a secondary counter called the 'denominator' in the schema. The client must use the provided and secondary counters to compute the final intended value.

For example:

```
Determining the average wait time for a workload per visit
Query the 'wait_time' and 'visits' field from a 'qos_detail' row:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/qos_detail/rows/<instance-
id>?fields=counters&counters.name=visits&#124;wait_time"
{
"counter_table": {
  "name": "qos_detail"
},
"id": "main-vsim1:_WAFL.CPU_ha",
"counters": [
  {
    "name": "visits",
    "value": 14631
  },
  {
    "name": "wait_time",
    "value": 167816
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/counter/tables/qos_detail/rows/<instance-id>"
  }
}
}
The average wait time per visit is calculated as 167816 / 14631 = 11
micro-seconds
```

In the above example, the average is calculated since boot-time. Sample periods are discussed in more detail below.

### Counter Computations

The statistics available through the counter tables gives you information about a specific point in time. This data can be useful, but more often you are interested in the statistics over a period of time.

The procedure for calculating a statistic over a period of time involves the following:

- Collect a data sample at the beginning of the period. If the counter requires a denominator, this should be collected at the same time.
- Collect a second data sample at the end of the period. If the counter requires a denominator, collect a second sample at the same time.
- Calculate the final result using the collected information and the formula associated with the counter type below

All counters that are not of type 'raw' will require some computation to be useful.

```
----------------

'''

T1, T2 : The start and ending time of the sample period
C1, C2 : The counter value at the start and ending time of the period
D1, D2 : The denominator value at the start and ending time of the period
----------------

'''

Percentage = ((C2 - C1) x 100) / (D2 - D1)
Rate = (C2 - C1) / (T2 - T1)
Average = (C2 - C1) / (D2 - D1)
Delta = C2 - C1
----------------

'''
```

### Aggregation

An aggregation is a logical container that consolidates the information from multiple entities into a single entity. There are two methods of aggregating tables:

- Automatic
- Combination.

### Automatic

Tables with automatic aggregation are generated by consolidating all entities with matching identifiers. The underlying tables that contribute to the aggregated table are referenced by the following syntax: {table_name}:constituent.

### Combination

Tables with combination aggregation are generated by consolidating all entities according to a unique field in the definition. The name of the combination table uses the following syntax: {table name}:{aggregation_name}.

An example combination table is 'volume:svm' table. This table aggregates all the volume statistics associated with a given vserver into a single table.

### Multi-Dimensional Arrays

Numeric counters can be scalar, one-dimensional or two dimensional values. Scalars are the most common values which consist of a single numeric value. A one-dimensional array is commonly used to present histograms such as the following table:

```
< 1s     :  3
< 5s     : 10
< 60s    :  1
```

A counter endpoint response that contains the above table would be formatted as follows:

```
{
"name": "Sample One-Dimensional Counter",
"labels": [ "< 1s", "< 5s", "< 60s" ],
"values": [3, 10, 1]
}
```

A two-dimensional array is used to report information about more complex relationships. An example data set is below:

```
            New      Used
--------------------------

'''

Car             1         2
Truck           3         4
Motorcycle      5         6
```

A counter endpoint response that contains the above table would be formatted as follows:

```
{
"name": "Sample Two-Dimensional Counter",
"labels": [ "New", "Used" ],
"counters": [
  {
    "label": "Car",
    "values": [1, 2]
  },
  {
    "label": "Truck",
    "values": [3, 4]
  },
  {
    "label": "Motorcycle",
    "values": [5, 6]
  }
]
}
```

#### Filtering / Querying

The counter endpoints adhere to the same behavior as other endpoints, with exception of how queries are handled for nested array fields.

The default behavior when processing a nested array query is to return the entire array content on a match. The counter endpoints' behavior will only return entries in the array that match the query.

Counter responses can contain a significant amount of data. This behavior improves the response by only returning the information requested and eliminating extra work for the client.

For example:

```
Given the following array:
"list": [ "fruit_apple", "color_red" ]
When you apply the following query:
list=fruit*
The default query behavior will return the array as:
"list": [ "fruit_apple", "color_red" ]
The counter endpoints will return the array as:
"list": [ "fruit_apple" ]
```

---

#### Examples

#### Retrieving a table schema definition

This example retrieves the table description and schema definition for the qos_detail table.

```
# The API:
/api/support/counter/tables/{name}

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/qos_detail?fields=*" -H "accept:
application/hal+json"

# The response:
{
"name": "qos_detail",
"description": "The qos_detail table that provides service center-based
statistical information.

*Note:*
This table returns a large number of rows. Querying by row name and using
wild cards may improve response times.",
"counter_schemas": [
  {
    "name": "in_latency_path",
    "description": "Determines whether or not service center-based
statistics are in the latency path.",
    "type": "raw",
    "unit": "none"
  },
  {
    "name": "node.name",
    "description": "System node name",
    "type": "string",
    "unit": "none"
  },
  {
    "name": "resource.name",
    "description": "Name of the associated resource.",
    "type": "string",
    "unit": "none"
  },
  {
    "name": "service_time",
    "description": "The workload's average service time per visit to the
service center.",
```

```
    "type": "average",
    "unit": "microsec",
    "denominator": {
      "name": "visits"
    }
  },
  {
    "name": "visits",
    "description": "The number of visits that the workload made to the
service center; measured in visits per second.",
    "type": "rate",
    "unit": "per_sec"
  },
  {
    "name": "wait_time",
    "description": "The workload's average wait time per visit to the
service center.",
    "type": "average",
    "unit": "microsec",
    "denominator": {
      "name": "visits"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/counter/tables/qos_detail"
  }
}
}
```


#### Query for tables that contain a keyword in the description

This example retrieves all table definitions contain the word "security" in their description.

```
# The API:
/api/support/counter/tables

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/?fields=name,description&description=*secur
ity*" -H "accept: application/hal+json"
```

```
# The response:
{
"records": [
  {
    "name": "csm_global",
    "description": "This table reports global statistics of the Cluster
Session Manager. The counters report the processing overhead of SpinNP
cryptography, both encryption and decryption, as carried out by CSM as it
handles cross-cluster data traffic, mostly on behalf of their data
protection operations. For example, a customer might seek to know the
processor time being consumed by these cryptographic operations in support
of their cross-cluster traffic. That data might help them evaluate the
performance impact of these security operations.",
    "_links": {
      "self": {
        "href": "/api/cluster/counter/tables/csm_global"
      }
    }
  },
  {
    "name": "file_directory",
    "description": "This table reports how many times file-directory jobs
were triggered to the set the file-security ACLS or SLAG ACLS. This
counter gives an indication how frequently the feature is being used to
set the ACLS on file-directory/volume.",
    "_links": {
      "self": {
        "href": "/api/cluster/counter/tables/file_directory"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/cluster/counter/tables?fields=name,description&description=*security
*"
  }
}
}
```

**Query for a specific property within all table rows.**

This example requests the property named 'node.name' for all 'wafl' table rows.

> The properties array content excludes any entries that do not match the provided query.

```
# The API:
/api/cluster/counter/tables/{counter_table.name}/rows

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/wafl/rows?properties.name=node.name&fields=
properties" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "id": "<instance id>",
    "properties": [
      {
        "name": "node.name",
        "value": "<node name>"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/cluster/counter/tables/wafl/rows/<instance id>"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/cluster/counter/tables/wafl/rows?properties.name=node.name&fields=pr
operties"
  }
}
}
```

**Query for a list of properties that match a wildcard on a specific row.**

This example queries for all properties associated with a row of the volume table.

> The properties array content excludes any entries that do not match the provided query.

```
# The API:
/api/cluster/counter/tables/{counter_table.name}/rows/{id}

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/volume/rows/<instance-
id>/?fields=properties&properties.name=svm*" -H "accept:
application/hal+json"

# The response:
{
"counter_table": {
  "name": "volume"
},
"id": "<instance-id>",
"properties": [
  {
    "name": "svm.name",
    "value": "<svm-name>"
  },
  {
    "name": "svm.uuid",
    "value": "4774d11c-a606-11ec-856f-005056bb7b59"
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/counter/tables/volume/rows/<instance-id>/"
  }
}
}
```

**Query for a list of counters in a specific table row**

This example queries for an explicit list of counters within a single row of the wafl table.

> The counters array content excludes any entries that do not match the provided query.

```
# The API:
/api/cluster/counter/tables/{counter_table.name}/rows/{id}

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/counter/tables/wafl/rows/<instance-
id>?fields=counters&counters.name=memory_used&#124;memory_free" -H
"accept: application/hal+json"

# The response:
{
"counter_table": {
  "name": "wafl"
},
"id": "<instance-id>",
"counters": [
  {
    "name": "memory_used",
    "value": 541
  },
  {
    "name": "memory_free",
    "value": 786
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/counter/tables/wafl/rows/<instance-id>"
  }
}
}
```

**Retrieve counter tables with schema definitions**

GET `/cluster/counter/tables`

**Introduced In:** 9.11

Returns a collection of counter tables and their schema definitions.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| counter_schemas.unit | string | query | False | Filter by counter_schemas.unit |
| counter_schemas.type | string | query | False | Filter by counter_schemas.type |
| counter_schemas.denominator.name | string | query | False | Filter by counter_schemas.denominator.name |
| counter_schemas.name | string | query | False | Filter by counter_schemas.name |
| counter_schemas.description | string | query | False | Filter by counter_schemas.description |
| description | string | query | False | Filter by description |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |
| Limit the number of records returned. | return_records | boolean | query | False |
| The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 | return_timeout | integer | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[counter_table] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "counter_schemas": [
        {
          "denominator": {
            "name": "string"
          },
          "description": "string",
          "name": "string",
          "type": "string",
          "unit": "string"
        }
      ],
      "description": "string",
      "name": "string"
    }
  ]
}
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 8585368 | The system has not completed it's initialization |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

counter_denominator

Counter used as the denominator in calculating the resulting value of averages and percentages.

| Name | Type | Description |
|---|---|---|
| name | string | Counter name. |

counter_schema

Schema definition of a single counter or property.

| Name | Type | Description |
|---|---|---|
| denominator | counter_denominator | Counter used as the denominator in calculating the resulting value of averages and percentages. |
| description | string | Counter or property description. |
| name | string | Counter or property name. |
| type | string | Type of counter or property. Properties will always set this field to 'string'. |
| unit | string | Counter unit. |

counter_table

Information for a single counter table.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| counter_schemas | array[counter_schema] | Array of counter schema definitions. |
| description | string | Description of the table. |
| name | string | Table name. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve counter rows**

GET `/cluster/counter/tables/{counter_table.name}/rows`

**Introduced In:** 9.11

Returns a collection of counter rows.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| counter_table.name | string | path | True | Counter table name. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| counters.value | integer | query | False | Filter by counters.value |
| counters.values | integer | query | False | Filter by counters.values |
| counters.counters.values | integer | query | False | Filter by counters.counters.values |
| counters.counters.label | string | query | False | Filter by counters.counters.label |
| counters.labels | string | query | False | Filter by counters.labels |
| counters.name | string | query | False | Filter by counters.name |
| id | string | query | False | Filter by id |
| properties.value | string | query | False | Filter by properties.value |
| properties.name | string | query | False | Filter by properties.name |
| aggregation.complete | boolean | query | False | Filter by aggregation.complete |
| aggregation.count | integer | query | False | Filter by aggregation.count |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| Limit the number of records returned. | return_records | boolean | query | False |
| The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 | return_timeout | integer | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[counter_row] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "counter_table": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "counters": [
        {
          "counters": [
            {
              "label": "string",
              "values": [
                "integer"
              ]
            }
          ],
          "labels": [
            "string"
          ],
          "name": "string",
          "values": [
            "integer"
          ]
        }
      ],
```

```
      "id": "string",
      "properties": [
        {
          "name": "string",
          "value": "string"
        }
      ]
    }
  ]
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585320 | Table requested is not found |
| 8586228 | Invalid counter name request. |
| 8586229 | Invalid counter property request. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

instance_counter_aggregation

Aggregation information about this counter.

| Name | Type | Description |
|---|---|---|
| complete | boolean | The aggregation state for this row.<br>For non-aggregated tables:<br>Not present<br>For aggregated tables:<br>If all requests to remote nodes for counter data are successful, then this value will be 'true'.<br>If any requests to remote nodes fail, then this value will be 'false'. |
| count | integer | Number of nodes included in the aggregation of this counter. |

counter_table_reference

Counter table reference.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Counter table name. |

counter2d

Counters that represent the second dimension of a two-dimension counter.

| Name | Type | Description |
| --- | --- | --- |
| label | string | Second dimension label. |
| values | array[integer] | List of values for the counter. |

counter

Representation of a counter and contains one of the following:

- Scalar counter populates the 'name' and 'value' fields.
- A 1D array populates the 'name', 'labels' and 'values' fields.
- A 2D array is represented as a list of counter entries.

```
"counters": [
  // Scalar counter
  {
      "name": "memory",
      "value": 4480
  },
  // one dimensional array "sys_read_latency_hist"
  {
      "name": "sys_read_latency_hist",
      "labels": ["0 - &lt;1ms", "1 - &lt;2ms", ...],
      "values": [0, 0, ...]
  },
  // Two dimensional array "foo" with ["Label 1", "Label 2"] as the
first
  // array dimension and labels ["w", "x", "y"] for the 2nd dimension
  {
      "name": "foo",
      "labels": ["Label 1", "Label 2"],
      "counters": [
          {
              "label": "x",
              "values": [0, 0]
          },
          {
              "label": "y",
              "values": [0, 0]
          },
          {
              "label": "z",
              "values": [0, 0]
          }
      ]
  }
```

| Name | Type | Description |
|---|---|---|
| counters | array[counter2d] | List of labels and values for the second dimension. |
| labels | array[string] | List of labels for the first dimension. |
| name | string | Counter name. |
| value | integer | Scalar value. |

| Name | Type | Description |
| --- | --- | --- |
| values | array[integer] | List of values in a one-dimensional counter. |

counter_property

Single string counter entry.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Property name. |
| value | string | Property value. |

counter_row

A single row of counter and property counter data.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| aggregation | instance_counter_aggregation | Aggregation information about this counter. |
| counter_table | counter_table_reference | Counter table reference. |
| counters | array[counter] | Array of counter name/value pairs. |
| id | string | Unique row identifier. |
| properties | array[counter_property] | Array of property name/value pairs. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a counter row

GET `/cluster/counter/tables/{counter_table.name}/rows/{id}`

**Introduced In:** 9.11

Returns a single counter row.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| counter_table.name | string | path | True | Counter table name. |
| id | string | path | True | Unique row identifier. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| aggregation | instance_counter_aggregation | Aggregation information about this counter. |
| counter_table | counter_table_reference | Counter table reference. |
| counters | array[counter] | Array of counter name/value pairs. |

| Name | Type | Description |
| --- | --- | --- |
| id | string | Unique row identifier. |
| properties | array[counter_property] | Array of property name/value pairs. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "counter_table": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string"
  },
  "counters": [
    {
      "counters": [
        {
          "label": "string",
          "values": [
            "integer"
          ]
        }
      ],
      "labels": [
        "string"
      ],
      "name": "string",
      "values": [
        "integer"
      ]
    }
  ],
  "id": "string",
  "properties": [
    {
      "name": "string",
      "value": "string"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585320 | Table requested is not found |
| 8586228 | Invalid counter name request. |
| 8586229 | Invalid counter property request. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

instance_counter_aggregation

Aggregation information about this counter.

| Name | Type | Description |
|---|---|---|
| complete | boolean | The aggregation state for this row.<br>For non-aggregated tables:<br>Not present<br>For aggregated tables:<br>If all requests to remote nodes for counter data are successful, then this value will be 'true'.<br>If any requests to remote nodes fail, then this value will be 'false'. |
| count | integer | Number of nodes included in the aggregation of this counter. |

counter_table_reference

Counter table reference.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Counter table name. |

counter2d

Counters that represent the second dimension of a two-dimension counter.

| Name | Type | Description |
|---|---|---|
| label | string | Second dimension label. |

| Name | Type | Description |
| --- | --- | --- |
| values | array[integer] | List of values for the counter. |

counter

Representation of a counter and contains one of the following:

- Scalar counter populates the 'name' and 'value' fields.
- A 1D array populates the 'name', 'labels' and 'values' fields.
- A 2D array is represented as a list of counter entries.

```
"counters": [
// Scalar counter
{
    "name": "memory",
    "value": 4480
},
// one dimensional array "sys_read_latency_hist"
{
    "name": "sys_read_latency_hist",
    "labels": ["0 - &lt;1ms", "1 - &lt;2ms", ...],
    "values": [0, 0, ...]
},
// Two dimensional array "foo" with ["Label 1", "Label 2"] as the first
// array dimension and labels ["w", "x", "y"] for the 2nd dimension
{
    "name": "foo",
    "labels": ["Label 1", "Label 2"],
    "counters": [
        {
            "label": "x",
            "values": [0, 0]
        },
        {
            "label": "y",
            "values": [0, 0]
        },
        {
            "label": "z",
            "values": [0, 0]
        }
    ]
}
```

| Name | Type | Description |
| --- | --- | --- |
| counters | array[counter2d] | List of labels and values for the second dimension. |
| labels | array[string] | List of labels for the first dimension. |
| name | string | Counter name. |
| value | integer | Scalar value. |
| values | array[integer] | List of values in a one-dimensional counter. |

counter_property

Single string counter entry.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Property name. |
| value | string | Property value. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve counter table details

GET `/cluster/counter/tables/{name}`

**Introduced In:** 9.11

Returns the information about a single counter table.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Counter table name. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| counter_schemas | array[counter_schema] | Array of counter schema definitions. |
| description | string | Description of the table. |
| name | string | Table name. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "counter_schemas": [
    {
      "denominator": {
        "name": "string"
      },
      "description": "string",
      "name": "string",
      "type": "string",
      "unit": "string"
    }
  ],
  "description": "string",
  "name": "string"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585320 | Table requested is not found |
| 8585368 | The system has not completed it's initialization |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

counter_denominator

Counter used as the denominator in calculating the resulting value of averages and percentages.

| Name | Type | Description |
|---|---|---|
| name | string | Counter name. |

counter_schema

Schema definition of a single counter or property.

| Name | Type | Description |
|---|---|---|
| denominator | counter_denominator | Counter used as the denominator in calculating the resulting value of averages and percentages. |
| description | string | Counter or property description. |
| name | string | Counter or property name. |
| type | string | Type of counter or property. Properties will always set this field to 'string'. |
| unit | string | Counter unit. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster firmware history

### Manage cluster firmware history

#### Overview

Use this API to retrieve a history of firmware update requests. This API supports GET calls.

---

#### Examples

### Retrieving history of firmware updates

The following example retrieves a history of firmware updates performed on the cluster.
Note that if the *fields=** parameter is not specified, only the job ID and start time are returned.
Filters can be added on the fields to limit the results.

```
# The API:
GET /api/cluster/firmware/history

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/firmware/history/?fields=*" -H
"accept: application/hal+json"

# The response:
200 OK
{
"records": [
  {
    "start_time": "1970-01-01T00:02:03+00:00",
    "job": {
      "uuid": "adf700c2-b50e-11ea-a54f-005056bbec43"
    },
    "node": {
```

```
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node1"
    },
    "fw_file_name": "all_disk_fw.zip",
    "fw_update_state": "starting_workers",
    "end_time": "1970-01-01T00:07:36+00:00",
    "update_status": [
      {
        "worker": {
          "node": {
            "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
            "name": "node1"
          },
          "state": "failed",
          "error": {
            "message": "A firmware file already exists.",
            "code": 2228327
          }
        }
      },
      {
        "worker": {
          "node": {
            "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ef",
            "name": "node2"
          },
          "state": "complete",
          "error": {
            "message": "Success",
            "code": 0
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/cluster/firmware/history/1970-01-01T00%3A02%3A03-
00%3A00/adf700c2-b50e-11ea-a54f-005056bbec43"
      }
    }
  },
  {
    "start_time": "1970-01-01T00:02:03+00:00",
    "job": {
      "uuid": "f84adabe-b50e-11ea-a54f-005056bbec43"
    },
```

```
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node1"
    },
    "fw_file_name": "all_shelf_fw.zip",
    "fw_update_state": "completed",
    "end_time": "1970-01-01T00:07:36+00:00",
    "update_status": [
      {
        "worker": {
          "node": {
            "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
            "name": "node1"
          },
          "state": "failed",
          "error": {
            "message": "A firmware file already exists.",
            "code": 2228327
          }
        }
      },
      {
        "worker": {
          "node": {
            "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ef",
            "name": "node2"
          },
          "state": "complete",
          "error": {
            "message": "Success",
            "code": 0
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/cluster/firmware/history/1970-01-01T00%3A02%3A03-
00%3A00/f84adabe-b50e-11ea-a54f-005056bbec43"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
```

```
      "href": "/api/cluster/firmware/history/?fields=%2A"
    }
  }
  }
```


---

### Retrieve history details for firmware update requests

GET `/cluster/firmware/history`

**Introduced In:** 9.8

Retrieves the history details for firmware update requests.

#### Learn more

- DOC /cluster/firmware/history

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| update_status.worker.node.name | string | query | False | Filter by update_status.worker.node.name |
| update_status.worker.node.uuid | string | query | False | Filter by update_status.worker.node.uuid |
| update_status.worker.error.code | integer | query | False | Filter by update_status.worker.error.code |
| update_status.worker.error.message | string | query | False | Filter by update_status.worker.error.message |
| update_status.worker.state | string | query | False | Filter by update_status.worker.state |
| fw_update_state | string | query | False | Filter by fw_update_state |
| end_time | string | query | False | Filter by end_time |
| job.uuid | string | query | False | Filter by job.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| start_time | string | query | False | Filter by start_time |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| fw_file_name | string | query | False | Filter by fw_file_name |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[firmware_history] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "end_time": "2019-02-02 19:00:00 +0000",
      "fw_file_name": "all_disk_fw.zip",
      "fw_update_state": "string",
      "job": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "string"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "start_time": "2019-02-02 19:00:00 +0000",
      "update_status": [
        {
          "worker": {
            "error": {
              "code": 2228325,
              "message": "Cannot open local staging ZIP file
```

```
disk_firmware.zip"
            },
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "state": "waiting_to_retry"
          }
        }
      ]
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

firmware_history_update_state_error

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the status message. |
| message | string | Error message returned when a firmware update job fails. |

worker

| Name | Type | Description |
| --- | --- | --- |
| error | firmware_history_update_state_error | |
| node | node | |
| state | string | The state of each worker that a node is controlling. |

firmware_history_update_state

| Name | Type | Description |
| --- | --- | --- |
| worker | worker | |

firmware_history

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| end_time | string | End time of this update request. |
| fw_file_name | string | Name of the firmware file. |
| fw_update_state | string | |
| job | job_link | |
| node | node | |
| start_time | string | Start time of this update request. |
| update_status | array[firmware_history_update_state] | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and manage cluster jobs

### View and manage cluster jobs

#### Overview

You can use this API to view and manipulate jobs. Jobs provide information about asynchronous operations. Some long-running jobs are paused or cancelled by calling a PATCH request. Individual operations indicate if they support PATCH requests on the job. After a job transitions to a terminal state, it is deleted after a default time of 300 seconds. Attempts to call a GET or PATCH request on the job returns a 404 error code After the job has been deleted.

#### Example

The following examples show how to retrieve and update a job state:

### Retrieving job information

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-
005056bbd8f5" -H "accept: application/json"

# The response:
{
  "uuid": "b5145e1d-b53b-11e8-8252-005056bbd8f5",
  "code": 0,
  "description": "Cluster Backup Job",
  "state": "running",
  "message": "creating_node_backups",
  "_links": {
      "self": {
          "href": "/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5"
      }
  }
}
```

### Updating a job that supports the new state

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-
005056bbd8f5?action=cancel" -H "accept: application/json"
```

---

### Retrieve recent asynchronous jobs

GET `/cluster/jobs`

**Introduced In:** 9.6

Retrieves a list of recently running asynchronous jobs. After a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| code | integer | query | False | Filter by code |
| end_time | string | query | False | Filter by end_time |
| message | string | query | False | Filter by message |
| uuid | string | query | False | Filter by uuid |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.8 |
| svm.uuid | string | query | False | Filter by svm.uuid<br>• Introduced in: 9.8 |
| description | string | query | False | Filter by description |
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| error.arguments.message | string | query | False | Filter by error.arguments.message<br>• Introduced in: 9.9 |
| error.arguments.code | string | query | False | Filter by error.arguments.code<br>• Introduced in: 9.9 |
| error.code | string | query | False | Filter by error.code<br>• Introduced in: 9.9 |
| error.message | string | query | False | Filter by error.message<br>• Introduced in: 9.9 |
| start_time | string | query | False | Filter by start_time |
| node.name | string | query | False | Filter by node.name<br>• Introduced in: 9.11 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[job] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "code": 0,
      "description": "App Snapshot Job",
      "end_time": "string",
      "error": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "message": "Complete: Successful",
      "node": {
        "name": "string"
      },
      "start_time": "string",
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

The error that caused the job to fail. This property is only populated when the job fails and it matches the API response error structure used by all APIs. The message and code match the dedicated message and code properties once the job has failed.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

node

The node where this job was run

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

job

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| code | integer | If the state indicates "failure", this is the final error code. |
| description | string | The description of the job to help identify it independent of the UUID. |
| end_time | string | The time the job ended. |
| error | error | The error that caused the job to fail. This property is only populated when the job fails and it matches the API response error structure used by all APIs. The message and code match the dedicated message and code properties once the job has failed. |
| message | string | A message corresponding to the state of the job providing additional details about the current state. |
| node | node | The node where this job was run |
| start_time | string | The time the job started. |
| state | string | The state of the job. |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve details of an asynchronous job**

GET `/cluster/jobs/{uuid}`

**Introduced In:** 9.6

Retrieves the details of a specific asynchronous job. After a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Job UUID |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| code | integer | If the state indicates "failure", this is the final error code. |
| description | string | The description of the job to help identify it independent of the UUID. |
| end_time | string | The time the job ended. |
| error | error | The error that caused the job to fail. This property is only populated when the job fails and it matches the API response error structure used by all APIs. The message and code match the dedicated message and code properties once the job has failed. |
| message | string | A message corresponding to the state of the job providing additional details about the current state. |
| node | node | The node where this job was run |
| start_time | string | The time the job started. |
| state | string | The state of the job. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "code": 0,
  "description": "App Snapshot Job",
  "end_time": "string",
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "message": "Complete: Successful",
  "node": {
    "name": "string"
  },
  "start_time": "string",
  "state": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

The error that caused the job to fail. This property is only populated when the job fails and it matches the API response error structure used by all APIs. The message and code match the dedicated message and code properties once the job has failed.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

node

The node where this job was run

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the state of an asynchronous job**

PATCH `/cluster/jobs/{uuid}`

**Introduced In:** 9.6

Updates the state of a specific asynchronous job.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Job UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| action | string | query | False | Requests a job to pause, resume, or cancel. Note that not all jobs support these actions. A job can only be resumed if it is in a paused state. After you successfully request a job to be cancelled, the job state changes to either success or failure.<br><br>• enum: ["pause", "resume", "cancel"] |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 458762 | Job is already in a terminal state. |
| 458773 | The Job Manager is not initialized. |
| 458771 | The specified job is running. |
| 458776 | The specified job is not currently running. |
| 458783 | This job does not support pause. |
| 458784 | This job does not support cancel. |

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve capacity pool licenses

### Retrieve capacity pool licenses

#### Overview

Capacity pool licenses are installed on and managed by the license manager. Each ONTAP node that is using the capacity pools licensing model is associated with a capacity pool license from which capacity is leased for data aggregates.

This API is used to retrieve information about associations between ONTAP nodes in the cluster and capacity pool licenses. It also reports how much capacity each node is consuming from the capacity pool.

---

#### Examples

### Retrieving a collection of capacity pools associated with the cluster

This example retrieves a collection that contains two capacity pool licenses, each of which is associated with an HA pair of nodes in a four-node cluster.

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/capacity-pools"
```

```
# Response
200 OK

# JSON Body
{
"records":[
  {
    "serial_number":"390000100",
    "license_manager": {
      "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-
11e0-ae1c-112233445566"
         }
      }
    },
    "nodes":[
      {
        "node":{
          "name":"node-1",
          "uuid":"4ea7a442-86d1-11e0-ae1c-123478563411"
        },
        "used_size":1099511627776,
        "_links":{
          "self":{
            "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-
123478563411"
          }
        }
      },
      {
        "node":{
          "name":"node-2",
          "uuid":"4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "used_size":1099511627776,
        "_links":{
          "self":{
            "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-
123478563412"
          }
        }
      }
    ],
    "_links":{
```

```
      "self":{
        "href":"/api/cluster/licensing/capacity-pools/390000100"
      }
    }
  },
  {
    "serial_number":"390000101",
    "license_manager": {
      "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-
11e0-ae1c-112233445566"
         }
      }
    },
    "nodes":[
      {
        "node":{
          "name":"node-3",
          "uuid":"4ea7a442-86d1-11e0-ae1c-123478563413"
        },
        "used_size":2199023255552,
        "_links":{
          "self":{
            "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-
123478563413"
          }
        }
      },
      {
        "node":{
          "name":"node-4",
          "uuid":"4ea7a442-86d1-11e0-ae1c-123478563414"
        },
        "used_size":2199023255552,
        "_links":{
          "self":{
            "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-
123478563414"
          }
        }
      }
    ],
    "_links":{
      "self":{
```

```
        "href":"/api/cluster/licensing/capacity-pools/390000101"
      }
    }
  }
],
"num_records":2,
"_links":{
  "self":{
    "href":"/api/cluster/licensing/capacity-pools"
  }
}
}
```


#### Retrieving information about nodes associated with a specific capacity pool license

This example retrieves information about the nodes that are associated with a capacity pool license of the serial number 390000100.


```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/capacity-
pools/390000100"

# Response
200 OK

# JSON Body
{
"serial_number":"390000100",
"license_manager": {
  "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-11e0-
ae1c-112233445566"
     }
  }
},
"nodes":[
  {
    "node":{
      "name":"node-1",
      "uuid":"4ea7a442-86d1-11e0-ae1c-123478563411"
    },
    "used_size":1099511627776,
    "_links":{
      "self":{
```

```
          "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563411"
        }
      }
    },
    {
      "node":{
        "name":"node-2",
        "uuid":"4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "used_size":1099511627776,
      "_links":{
        "self":{
          "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      }
    }
  ],
  "_links":{
    "self":{
      "href":"/api/cluster/licensing/capacity-pools/390000100"
    }
  }
  }
```


**Retrieve capacity pools**

GET `/cluster/licensing/capacity-pools`

**Introduced In:** 9.8

Retrieves a collection of capacity pools.

#### Learn more

- DOC /cluster/licensing/capacity-pools

#### Related ONTAP commands

- `system license show-status`
- `system license show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nodes.used_size | integer | query | False | Filter by nodes.used_size |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nodes.node.name | string | query | False | Filter by nodes.node.name |
| nodes.node.uuid | string | query | False | Filter by nodes.node.uuid |
| serial_number | string | query | False | Filter by serial_number |
| license_manager.uuid | string | query | False | Filter by license_manager.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "license_manager": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
      },
      "nodes": [
        {
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "used_size": 0
        }
      ],
      "serial_number": 390000100
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

license_manager

License manager instance where this capacity pool license in installed.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| uuid | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

nodes

Information on a node from the capacity licensing perspective.

| Name | Type | Description |
|---|---|---|
| node | node_reference | |

| Name | Type | Description |
|---|---|---|
| used_size | integer | Capacity, in bytes, that is currently used by the node. |

records

Information on a capacity pool license and how it is associated with the cluster.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| license_manager | license_manager | License manager instance where this capacity pool license in installed. |
| nodes | array[nodes] | Nodes in the cluster associated with this capacity pool. |
| serial_number | string | Serial number of the capacity pool license. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve capacity pool information

GET `/cluster/licensing/capacity-pools/{serial_number}`

**Introduced In:** 9.8

Retrieves information about the capacity pool.

#### Learn more

- DOC /cluster/licensing/capacity-pools

#### Related ONTAP commands

- `system license show-status`
- `system license show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| serial_number | string | path | True | Serial number of the capacity pool license. |
| nodes.used_size | integer | query | False | Filter by nodes.used_size |
| nodes.node.name | string | query | False | Filter by nodes.node.name |
| nodes.node.uuid | string | query | False | Filter by nodes.node.uuid |
| license_manager.uuid | string | query | False | Filter by license_manager.uuid |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| license_manager | license_manager | License manager instance where this capacity pool license in installed. |

| Name | Type | Description |
| --- | --- | --- |
| nodes | array[nodes] | Nodes in the cluster associated with this capacity pool. |
| serial_number | string | Serial number of the capacity pool license. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "license_manager": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
  },
  "nodes": [
    {
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "used_size": 0
    }
  ],
  "serial_number": 390000100
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

license_manager

License manager instance where this capacity pool license in installed.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| uuid | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

nodes

Information on a node from the capacity licensing perspective.

| Name | Type | Description |
|---|---|---|
| node | node_reference | |
| used_size | integer | Capacity, in bytes, that is currently used by the node. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster license managers

### Manage cluster license managers

#### Overview

This API is used to manage information about the license manager instance associated with the cluster.

When an ONTAP cluster is initially created to use the capacity pools licensing model, information about the license manager instance that the cluster should use is pre-configured. Generally, this configuration does not need to be updated unless the license manager instance changes its IP address.

The license manager is currently bundled with the ONTAP Select Deploy utility and runs on the same VM as ONTAP Select Deploy. Use this API to update the license manager IP address when the Deploy VM changes its IP address.

---

#### Examples

### Retrieving information about the license manager instance associated with the cluster

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/license-managers"

# Response
200 OK

# JSON Body
{
"records": [
  {
    "uuid":"4ea7a442-86d1-11e0-ae1c-112233445566",
    "uri": {
      "host":"10.1.1.1",
    },
    "default":true
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/cluster/licensing/license-managers"
  }
}
}
```

#### Updating an existing license manager instance

```
# API
curl -X PATCH "https://<mgmt-ip>/api/cluster/licensing/license-
managers/4ea7a442-86d1-11e0-ae1c-112233445566"

# JSON Body
{
"uri": {
  "host":"10.1.1.3"
}
}

# Response
202 Accepted
```

#### Retrieve license managers

GET `/cluster/licensing/license-managers`

**Introduced In:** 9.8

Retrieves a collection of license managers.

#### Learn more

- DOC /cluster/licensing/license-managers

#### Related ONTAP commands

- `system license license-manager show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uri.host | string | query | False | Filter by uri.host |
| default | boolean | query | False | Filter by default |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uri": {
        "host": "10.1.1.1"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

uri

License manager URI.

| Name | Type | Description |
|---|---|---|
| host | string | License manager host name, IPv4 or IPv6 address. |

records

Information on a license manager instance associated with the cluster.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| default | boolean | Flag that indicates whether it's the default license manager instance used by the cluster.' When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance. |
| uri | uri | License manager URI. |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the license manager information

GET `/cluster/licensing/license-managers/{uuid}`

**Introduced In:** 9.8

Retrieves information about the license manager.

#### Learn more

- DOC /cluster/licensing/license-managers

#### Related ONTAP commands

- `system license license-manager show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | |
| uri.host | string | query | False | Filter by uri.host |
| default | boolean | query | False | Filter by default |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| default | boolean | Flag that indicates whether it's the default license manager instance used by the cluster.'<br>When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance. |
| uri | uri | License manager URI. |
| uuid | string | |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uri": {
    "host": "10.1.1.1"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

uri

License manager URI.

| Name | Type | Description |
|---|---|---|
| host | string | License manager host name, IPv4 or IPv6 address. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the license manager configuration**

PATCH `/cluster/licensing/license-managers/{uuid}`

**Introduced In:** 9.8

Updates the license manager configuration.

#### Learn more

- DOC /cluster/licensing/license-managers

#### Related ONTAP commands

- `system license license-manager modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| default | boolean | Flag that indicates whether it's the default license manager instance used by the cluster.'<br>When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance. |
| uri | uri | License manager URI. |
| uuid | string | |

#### Example request

```
{
  "uri": {
    "host": "10.1.1.1"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115532 | The requested update to the license manager information failed. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| errors | array[returned_error] | |

#### Example error


```
{
  "errors": [
    {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  ]
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

uri

License manager URI.

| Name | Type | Description |
|---|---|---|
| host | string | License manager host name, IPv4 or IPv6 address. |

license_manager

Information on a license manager instance associated with the cluster.

| Name | Type | Description |
|---|---|---|
| default | boolean | Flag that indicates whether it's the default license manager instance used by the cluster.' When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance. |
| uri | uri | License manager URI. |
| uuid | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster licensing

### Manage cluster licensing

#### Overview

Licensing allows you to tailor a system to meet an organization's specific needs. You can enable new features by purchasing a license from a NetApp sales associate. After installation of the license, the new feature is available immediately.

This interface manages licenses according to their supported feature. By default, the interface displays packages with installed licenses, but you can also return unlicensed packages.

Each feature has a compliance state that is indicated at the package level. Individual licenses also contain a compliance state indicated in the "licenses" array. The state of the package is determined by analyzing the underlying licenses according to the following criteria:

- Licensing terms
- Cluster state

### Licensing terms

The licensing terms define the conditions under which a package is considered "compliant". Individual licenses are evaluated based on the following:

- Scope
- Time period
- Usage

### Scope

A package can be licensed under the following scopes:

- Site - Permits the feature to be used by any node in any cluster.
- Cluster - Permits the feature to be used by any node in a single specific cluster.
- Node - Permits the authorized node to use the feature. Within a cluster, if you don't supply every node with a valid license, the package state indicates "noncompliant". You must purchase a license for each node in a cluster for the package to be considered "compliant".

### Time period

Some package licenses are only valid for a limited period of time. After a license has expired, the package state changes to "noncompliant". You need to purchase a new license for the package to return to a "compliant" state.

### Usage

Some package licenses have additional terms that need to be maintained to keep a license in compliance. These conditions are defined by the individual license. For example, a license might define the maximum amount of storage that a node can allocate for the license to be "compliant".

### Cluster state

A cluster's state consists of the following:

- Node online status
- Node cluster membership

Some features require that a node be online to display a valid compliance state. If a node cannot be reached or is not known to the cluster, the individual license might indicate an "unknown" state.

---

#### Licensing keys

A license is issued in one of the following three formats:

- 28-character key
- NetApp License File Version 1 (NLFv1)
- NetApp License File Version 2 (NLFv2)

### Overview of NLFv1 and NLFv2 License Formats

NLFv1 and NLFv2 licenses are both JSON based files that allow features to be enabled.

The difference between the two formats is that a NLFv2 license allows multiple features to be enabled with a single file. A NLFv1 license is capable of enabling a single feature.

These licenses are identified, in the various methods, as follows:

| Format | Identifying Keys |
| --- | --- |
| 28 Character Key | name / serial_number |
| NLFv1 | name / serial_number |
| NLFv2 | licenses.installed_license / serial_number |

The following is an example of a 28-character key:

```
AMEPOSOIKLKGEEEEDGNDEKSJDEEE
```

The following is an example of an NLFv1 key:


```
{
"statusResp": {
"version": "1",
"serialNumber": "123456789",
"message": "Success",
"licenses": {
  "capacity": "1",
  "type": "capacity",
  "licenseProtocol": "FABRICPOOL-TB",
  "package": "FabricPool",
  "licenseScope": "cluster"
},
"snStatus": "Active",
"product": "fabricpool",
"statusCode": "S007"
},
"Signature": "signatureABC"
}
```


The following is an example of an NLFv2 key:

```
{
"statusResp": {
"version": "2",
"serialNumber": "123456789",
"message": "Success",
"product": "Sample NLFv2 License",
"licenses": {
  "capacity": "1",
  "type": "capacity",
  "HostID": "5554444",
  "package": [ "NFS", "CIFS" ],
  "licenseScope": "node"
},
"snStatus": "Active",
"statusCode": "S007"
},
"Signature": "signatureABC"
}
```


You can use this API to submit any format to enable features.

---

#### Examples

### Retrieving a collection of licenses organized by package

This example retrieves a collection that contains one entry for each package (filtered to only the 'fabricpool' package).

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?fields=*&name=fabricpool"

# Response
200 OK

# JSON Body
{
"records": [
{
  "name": "fabricpool",
  "scope": "cluster",
  "state": "compliant",
  "description": "FabricPool License",
  "licenses": [
  {
    "owner": "testcluster-1",
    "serial_number": "4149027342",
    "state": "compliant",
    "capacity": {
    "maximum_size": 1099511627776,
    "used_size": 0,
    "measurement_unit": "bytes"
    }
  }
  ],
  "_links": {
  "self": {
    "href": "/api/cluster/licensing/licenses/fabricpool"
  }
  }
}
],
"num_records": 1,
"_links": {
"self": {
  "href": "/api/cluster/licensing/licenses/?fields=*&name=fabricpool"
}
}
}
```

**Retrieving a collection of licenses organized by package - for package cloud**

The following example retrieves a collection that contains one entry for each package (filtered to only the

'cloud' package). The cloud package, in this example, is in the enforcement period as the license has expired. The REST GET output displays an additional field 'shutdown_imminent' to indicate that the system will shutdown.

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?fields=*&name=cloud"

# Response
200 OK

# JSON Body
{
"records": [
  {
    "name": "cloud",
    "scope": "node",
    "state": "noncompliant",
    "description": "Cloud ONTAP License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "unlicensed"
    },
    "licenses": [
      {
        "owner": "test-vsim1",
        "serial_number": "90120130000000000001",
        "active": false,
        "evaluation": true,
        "expiry_time": "2021-10-26T19:57:41Z",
        "shutdown_imminent": true,
        "compliance": {
          "state": "noncompliant"
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/cluster/licensing/licenses/cloud"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
```

```
    "href": "/api/cluster/licensing/licenses/?fields=*&name=cloud"
  }
}
}
```

#### Retrieving a collection of licenses installed with NLFv2

This example retrieves a collection of licenses that were installed by a NLFv2 formatted license.

> The license is referenced by the installed license "Core*Bundle" and the license serial number "4212426890"

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?fields=*&licenses.installed_license=Cor
e*Bundle&serial_number=4212426890"

# Response
200 OK

# JSON Body
{
"records": [
  {
    "name": "nfs",
    "scope": "node",
    "state": "noncompliant",
    "description": "NFS License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
```

```
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/nfs/?licenses.installed_license=Core*Bund
le"
      }
    }
  },
  {
    "name": "cifs",
    "scope": "node",
    "state": "noncompliant",
    "description": "CIFS License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
```

```
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/cifs/?licenses.installed_license=Core*Bun
dle"
      }
    }
  },
  {
    "name": "iscsi",
    "scope": "node",
    "state": "noncompliant",
    "description": "iSCSI License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
```

```
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/iscsi/?licenses.installed_license=Core*Bu
ndle"
      }
    }
  },
  {
    "name": "fcp",
    "scope": "node",
    "state": "noncompliant",
    "description": "FCP License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
```

```
      "self": {
        "href":
"/api/cluster/licensing/licenses/fcp/?licenses.installed_license=Core*Bund
le"
      }
    }
  },
  {
    "name": "snaprestore",
    "scope": "node",
    "state": "noncompliant",
    "description": "SnapRestore License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/snaprestore/?licenses.installed_license=C
ore*Bundle"
```

```
      }
    }
  },
  {
    "name": "flexclone",
    "scope": "node",
    "state": "noncompliant",
    "description": "FlexClone License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
      {
        "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/flexclone/?licenses.installed_license=Core*Bundle"
      }
    }
  },
  {
```

```
    "name": "nvme_of",
    "scope": "node",
    "state": "noncompliant",
    "description": "NVMe-oF License",
     "licenses": [
      {
         "owner": "test-vsim3",
         "active": false,
         "evaluation": false,
         "compliance": {
           "state": "unlicensed"
         }
       },
       {
         "owner": "test-vsim4",
         "installed_license": "Core Bundle",
         "host_id": "4212426890",
         "serial_number": "4212426890",
         "active": true,
         "evaluation": false,
         "compliance": {
           "state": "compliant"
         },
         "capacity": {
           "maximum_size": 10995116277760,
           "measurement_unit": "bytes"
         }
       }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/nvme_of/?licenses.installed_license=Core*
Bundle"
      }
    }
  },
  {
    "name": "s3",
    "scope": "node",
    "state": "noncompliant",
    "description": "S3 License",
    "entitlement": {
      "action": "acquire_license",
      "risk": "medium"
    },
```

```
    "licenses": [
      {
        "owner": "test-vsim3",
        "active": false,
        "evaluation": false,
        "compliance": {
          "state": "unlicensed"
        }
      },
     {
       "owner": "test-vsim4",
        "installed_license": "Core Bundle",
        "host_id": "4212426890",
        "serial_number": "4212426890",
        "active": true,
        "evaluation": false,
        "compliance": {
          "state": "compliant"
        },
        "capacity": {
          "maximum_size": 10995116277760,
          "measurement_unit": "bytes"
        }
      }
    ],
    "_links": {
      "self": {
        "href":
"/api/cluster/licensing/licenses/s3/?licenses.installed_license=Core*Bundl
e"
      }
    }
  }
],
"num_records": 8,
"_links": {
  "self": {
    "href":
"/api/cluster/licensing/licenses?fields=*&licenses.installed_license=Core*
Bundle&serial_number=4212426890"
  }
}
}
```

#### Retrieving a collection of installed licenses

This example retrieves a collection containing all packages (except base) that have installed licenses.

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?fields=*&name=!base"

# Response
200 OK

# JSON Body
{
"records": [
{
  "name": "nfs",
  "scope": "node",
  "state": "compliant",
  "description": "NFS License",
  "entitlement": {
    "action": "none",
    "risk": "low"
  },
  "licenses": [
  {
    "owner": "testcluster-1",
    "serial_number": "1-81-0000000000000004149027492",
    "state": "compliant"
  }
  ],
  "_links": {
  "self": {
    "href": "/api/cluster/licensing/licenses/nfs"
  }
  }
},
{
  "name": "cifs",
  "scope": "node",
  "state": "compliant",
  "description": "CIFS License",
  "entitlement": {
    "action": "acquire_license",
    "risk": "medium"
  },
  "licenses": [
```

```
    {
      "owner": "testcluster-1",
      "serial_number": "1-81-0000000000000004149027492",
      "state": "compliant"
    }
    ],
    "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/cifs"
    }
    }
  }
  ],
  "num_records": 2,
  "_links": {
  "self": {
    "href": "/api/cluster/licensing/licenses/?fields=*&name=!base"
  }
  }
  }
```


#### Retrieving a collection of unlicensed packages

By default, unlicensed packages are filtered from the collection output. This example shows how to use a query to retrieve unlicensed packages.

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?name=flexcache&state=unlicensed"

# Response
200 OK

# JSON Body
{
"records": [
  {
    "name": "flexcache",
    "_links": {
      "self": {
        "href": "/api/cluster/licensing/licenses/flexcache"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/cluster/licensing/licenses?name=flexcache&state=unlicensed"
  }
}
}
```

#### Installing a NLF license

This example installs a single NLFv1 license. A NLFv2 license installs using the same procedure.

> You must escape all the double quotes and backslash characters of the JSON license before it can be placed in the POST request.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses"

# JSON Body
{
"keys" : [ "{\"statusResp\":{\"snStatus\": \"Active\", \"licenses\": {
\"package\": \"FabricPool\", \"capacity\": \"1\", \"licenseProtocol\":
\"FABRICPOOL-TB\", \"type\": \"capacity\", \"licenseScope\": \"cluster\"},
\"message\": \"Success\", \"statusCode\": \"S007\", \"version\": \"1\",
\"product\": \"fabricpool\", \"serialNumber\": \"4149027342\"},
\"Signature\":\"SignatureABC\"}" ]
}

# Response
201 Created
```

#### Installing a 28-character key

This example installs a single 28-character key formatted license.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses"

# JSON Body
{
"keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA" ]
}

# Response
201 Created
```

#### Installing multiple licenses with one API call

This example shows how multiple keys can be provided to install multiple features in a single API call.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses"

# JSON Body
{
"keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA",
           "BBBBBBBBBBBBBBBBBBBBBBBBBBBB" ]
}

# Response
201 Created
```

#### Retrieving information for a specific license package

This example shows how to retrieve information about the specific feature package `fabricpool`.

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/licenses/fabricpool"

# Response
200 OK

# JSON Body
{
"name": "fabricpool",
"scope": "cluster",
"state": "compliant",
"description": "FabricPool License",
"licenses": [
{
  "owner": "testcluster-1",
  "serial_number": "123456789",
  "state": "compliant",
  "capacity": {
  "maximum_size": 109951162777600,
  "used_size": 0,
  "measurement_unit": "bytes"
  }
}
],
"_links": {
"self": {
  "href": "/api/cluster/licensing/licenses/fabricpool/"
}
}
}
```

#### Deleting a specific license

This example show how to delete a CIFS site license.

```
# API
curl -X DELETE "https://<mgmt-
ip>/api/cluster/licensing/licenses/cifs?serial_number=1-80-000011"

# JSON Body
{}

# Response
200 OK
```

#### Deleting with a query

The following example shows how to delete all NFS licenses specified with the '*' query.

```
# API
curl -X DELETE "https://<mgmt-
ip>/api/cluster/licensing/licenses/nfs?serial_number=*"

# JSON Body
{}

# Response
200 OK
```

#### Deleting all licenses installed with NLFv2

The following example shows how to delete all licenses installed by a NLFv2 formatted license.

```
# API
curl -X DELETE "https://<mgmt-
ip>/api/cluster/licensing/licenses?licenses.installed_license=Core*Bundle&
serial_number=4149026-97-8"

# JSON Body
{
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/cluster/licensing/licenses?licenses.installed_license=Core*Bundle&se
rial_number=4149026-97-8"
  }
}
}

# Response
200 OK
```

#### Retrieve license packages

GET `/cluster/licensing/licenses`

**Introduced In:** 9.6

Retrieves a collection of license packages.

By default, the GET method only returns licensed packages. You must provide the following query "state=unlicensed" to retrieve unlicensed packages.
**Note:** Starting with ONTAP 9.11.1, the GET method no longer returns the Base license record.

**Related ONTAP commands**

- `system license show-status`
- `system license show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| entitlement.risk | string | query | False | Filter by entitlement.risk<br>• Introduced in: 9.11 |
| entitlement.action | string | query | False | Filter by entitlement.action<br>• Introduced in: 9.11 |
| scope | string | query | False | Filter by scope |
| description | string | query | False | Filter by description<br>• Introduced in: 9.11 |
| state | string | query | False | Filter by state |
| name | string | query | False | Filter by name |
| licenses.owner | string | query | False | Filter by licenses.owner |
| licenses.host_id | string | query | False | Filter by licenses.host_id<br>• Introduced in: 9.9 |
| licenses.active | boolean | query | False | Filter by licenses.active |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| licenses.expiry_time | string | query | False | Filter by licenses.expiry_time |
| licenses.capacity.measurement_unit | string | query | False | Filter by licenses.capacity.measurement_unit<br>• Introduced in: 9.18 |
| licenses.capacity.maximum_size | integer | query | False | Filter by licenses.capacity.maximum_size |
| licenses.capacity.disabled_size | integer | query | False | Filter by licenses.capacity.disabled_size<br>• Introduced in: 9.18 |
| licenses.capacity.used_size | integer | query | False | Filter by licenses.capacity.used_size |
| licenses.serial_number | string | query | False | Filter by licenses.serial_number |
| licenses.evaluation | boolean | query | False | Filter by licenses.evaluation |
| licenses.shutdown_imminent | boolean | query | False | Filter by licenses.shutdown_imminent<br>• Introduced in: 9.11 |
| licenses.start_time | string | query | False | Filter by licenses.start_time |
| licenses.compliance.state | string | query | False | Filter by licenses.compliance.state |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| licenses.installed_license | string | query | False | Filter by licenses.installed_license<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "NFS License",
      "entitlement": {
        "action": "string",
        "risk": "string"
      },
      "licenses": [
        {
          "capacity": {
            "disabled_size": 0,
            "maximum_size": 0,
            "measurement_unit": "string",
            "used_size": 0
          },
          "compliance": {
            "state": "compliant"
          },
          "expiry_time": "2019-03-02 19:00:00 +0000",
          "host_id": "456-44-1234",
          "installed_license": "Core Bundle",
          "owner": "cluster1",
          "serial_number": "123456789",
          "start_time": "2019-02-02 19:00:00 +0000"
        }
      ],
      "name": "NFS",
      "scope": "string",
      "state": "compliant"
```

```
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

entitlement

| Name | Type | Description |
|---|---|---|
| action | string | Entitlement action to be taken to mitigate the risk |
| risk | string | Entitlement risk of the package |

capacity

| Name | Type | Description |
|---|---|---|
| disabled_size | integer | Capacity that is currently disabled due to license limits. Its unit of measure is specified in the measurement_unit field. |
| maximum_size | integer | Licensed capacity size that can be used. Its unit of measure is specified in the measurement_unit field. |
| measurement_unit | string | Unit of measure for capacity based licenses. |

| Name | Type | Description |
| --- | --- | --- |
| used_size | integer | Specifies the total number of GPUs in the system when measurement_unit is GPUs, else specifies the bytes used. |

compliance

| Name | Type | Description |
| --- | --- | --- |
| state | string | Compliance state of the license. |

licenses

| Name | Type | Description |
| --- | --- | --- |
| active | boolean | A flag indicating whether the license is currently being enforced. |
| capacity | capacity | |
| compliance | compliance | |
| evaluation | boolean | A flag indicating whether the license is in evaluation mode. |
| expiry_time | string | Date and time when the license expires. |
| host_id | string | A string that associates the license with a node or cluster. |
| installed_license | string | Name of license that enabled the feature. |
| owner | string | Cluster, node or license manager that owns the license. |
| serial_number | string | Serial number of the license. |

| Name | Type | Description |
| --- | --- | --- |
| shutdown_imminent | boolean | A flag indicating whether the Cloud ONTAP system is going to shutdown as the Cloud platform license has already expired.<br>• readOnly: 1<br>• Introduced in: 9.11<br>• x-nullable: true |
| start_time | string | Date and time when the license starts. |

records

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| description | string | License description |
| entitlement | entitlement | |
| licenses | array[licenses] | Installed licenses of the package. |
| name | string | Name of the license. |
| scope | string | Scope of the license. |
| state | string | Summary state of package based on all installed licenses. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Install one or more feature licenses

POST `/cluster/licensing/licenses`

**Introduced In:** 9.6

Installs one or more feature licenses.

#### Required properties

- `keys` - Array containing a list of NLF or 28-character license keys.

#### Related ONTAP commands

- `system license add`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| description | string | License description |
| entitlement | entitlement | |
| keys | array[string] | |
| licenses | array[licenses] | Installed licenses of the package. |
| name | string | Name of the license. |
| scope | string | Scope of the license. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Summary state of package based on all installed licenses. |

**Example request**

```
{
  "description": "NFS License",
  "entitlement": {
    "action": "string",
    "risk": "string"
  },
  "keys": [
    "<LICENSE-KEY>"
  ],
  "licenses": [
    {
      "capacity": {
        "disabled_size": 0,
        "maximum_size": 0,
        "measurement_unit": "string",
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 +0000",
      "host_id": "456-44-1234",
      "installed_license": "Core Bundle",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 +0000"
    }
  ],
  "name": "NFS",
  "scope": "string",
  "state": "compliant"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "description": "NFS License",
      "entitlement": {
        "action": "string",
        "risk": "string"
      },
      "keys": [
        "<LICENSE-KEY>"
      ],
      "licenses": [
        {
          "capacity": {
            "disabled_size": 0,
            "maximum_size": 0,
            "measurement_unit": "string",
            "used_size": 0
          },
          "compliance": {
            "state": "compliant"
          },
          "expiry_time": "2019-03-02 19:00:00 +0000",
          "host_id": "456-44-1234",
          "installed_license": "Core Bundle",
          "owner": "cluster1",
          "serial_number": "123456789",
          "start_time": "2019-02-02 19:00:00 +0000"
        }
      ],
      "name": "NFS",
      "scope": "string",
      "state": "compliant"
    }
  ]
}
```


### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1115117 | Generic licensing error |
| 1115122 | No cluster serial number found |
| 1115124 | No node serial number found |
| 1115130 | No license code was provided |
| 1115131 | Installation of the license failed |
| 1115132 | License already exists on system |
| 1115134 | Serial number does not belong to node |
| 1115141 | License data is invalid |
| 1115142 | License signature is invalid |
| 1115143 | Internal error applying the requested license |
| 1115152 | License does not apply to the platform |
| 1115154 | Unable to retrieve cluster ID |
| 1115155 | Invalid cluster ID found |
| 1115159 | License is not in an acceptable format |
| 1115160 | License has already expired |
| 1115164 | Minimum ONTAP version requirements not met |
| 1115165 | Minimum ONTAP version requirements are not met for license type enabled |
| 1115166 | Minimum ONTAP version requirements are not met for license protocol SEC-COMP-BNDL-ENBLD |
| 1115179 | FlexCache is not supported on this system |
| 1115180 | FlexCache is not supported on cloud systems |
| 1115407 | Capacity pool licenses cannot be installed directly |
| 1115427 | License is incompatible with capacity pools licensing mode |

| Error Code | Description |
|---|---|
| 1115562 | One or more errors occurred when installing a NLFv2 license |
| 1115563 | Package details and serial number of license contained within the NLFv2 failure |
| 1115564 | Package cannot be deleted individually as it is part of a bundle |
| 1115565 | NLFv2 install failed as the license serial number is already in use |
| 1115616 | Package details and serial number of license included in the install conflict |
| 1115617 | NLFv2 license install failed with summary of conflicting licenses |
| 1115618 | NLFv2 license install failed as a license with newer timestamp already exists |
| 5375355 | The cluster has more nodes than are supported by All SAN Array. |
| 5375366 | The cluster has one or more nodes that do not support All SAN Array. |
| 655294464 | Failed to extract license contents |
| 655294465 | License key is invalid |
| 655294466 | Serial number is invalid |
| 655294467 | Version number is invalid |
| 655294468 | Expired license |
| 655294469 | License does not apply to the platform |
| 655294470 | License does not apply to the product |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| errors | array[returned_error] | |

**Example error**

```
{
  "errors": [
    {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  ]
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

entitlement

| Name | Type | Description |
|---|---|---|
| action | string | Entitlement action to be taken to mitigate the risk |
| risk | string | Entitlement risk of the package |

capacity

| Name | Type | Description |
|---|---|---|
| disabled_size | integer | Capacity that is currently disabled due to license limits. Its unit of measure is specified in the measurement_unit field. |
| maximum_size | integer | Licensed capacity size that can be used. Its unit of measure is specified in the measurement_unit field. |
| measurement_unit | string | Unit of measure for capacity based licenses. |
| used_size | integer | Specifies the total number of GPUs in the system when measurement_unit is GPUs, else specifies the bytes used. |

compliance

| Name | Type | Description |
|---|---|---|
| state | string | Compliance state of the license. |

licenses

| Name | Type | Description |
| --- | --- | --- |
| active | boolean | A flag indicating whether the license is currently being enforced. |
| capacity | capacity | |
| compliance | compliance | |
| evaluation | boolean | A flag indicating whether the license is in evaluation mode. |
| expiry_time | string | Date and time when the license expires. |
| host_id | string | A string that associates the license with a node or cluster. |
| installed_license | string | Name of license that enabled the feature. |
| owner | string | Cluster, node or license manager that owns the license. |
| serial_number | string | Serial number of the license. |
| shutdown_imminent | boolean | A flag indicating whether the Cloud ONTAP system is going to shutdown as the Cloud platform license has already expired.<br>• readOnly: 1<br>• Introduced in: 9.11<br>• x-nullable: true |
| start_time | string | Date and time when the license starts. |

license_package

| Name | Type | Description |
| --- | --- | --- |
| description | string | License description |
| entitlement | entitlement | |
| keys | array[string] | |
| licenses | array[licenses] | Installed licenses of the package. |

| Name | Type | Description |
|---|---|---|
| name | string | Name of the license. |
| scope | string | Scope of the license. |
| state | string | Summary state of package based on all installed licenses. |

records

| Name | Type | Description |
|---|---|---|
| description | string | License description |
| entitlement | entitlement | |
| keys | array[string] | |
| licenses | array[licenses] | Installed licenses of the package. |
| name | string | Name of the license. |
| scope | string | Scope of the license. |
| state | string | Summary state of package based on all installed licenses. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a license

DELETE `/cluster/licensing/licenses/{name}`

**Introduced In:** 9.6

Deletes a license.

#### Related ONTAP commands

- `system license delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Name of the license package to delete. |
| serial_number | string | query | True | Serial number of the license to delete. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 525028 | Error during volume limit check, cannot remove license |
| 525029 | Current volume use will exceed limits if license is removed |
| 1115137 | Cluster license requires a base license to be installed |
| 1115144 | Cloud licenses cannot be deleted |
| 1115178 | A tier license that is still in use cannot be deleted |
| 1115213 | License is still in use and cannot be removed |
| 1115406 | Capacity pool licenses cannot be deleted |
| 1115564 | Package is part of a NLFv2 license and cannot be removed individually |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a license package

GET `/cluster/licensing/licenses/{name}`

**Introduced In:** 9.6

Retrieves a specific license package.

> By default, the GET method only returns licensed packages. You must provide the following query "state=unlicensed" to retrieve unlicensed packages.

#### Related ONTAP commands

- `system license show`
- `system license show-status`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Name of the license package. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| entitlement.risk | string | query | False | Filter by entitlement.risk<br><br>• Introduced in: 9.11 |
| entitlement.action | string | query | False | Filter by entitlement.action<br><br>• Introduced in: 9.11 |
| scope | string | query | False | Filter by scope |
| description | string | query | False | Filter by description<br><br>• Introduced in: 9.11 |
| state | string | query | False | Filter by state |
| licenses.owner | string | query | False | Filter by licenses.owner |
| licenses.host_id | string | query | False | Filter by licenses.host_id<br><br>• Introduced in: 9.9 |
| licenses.active | boolean | query | False | Filter by licenses.active |
| licenses.expiry_time | string | query | False | Filter by licenses.expiry_time |
| licenses.capacity.measurement_unit | string | query | False | Filter by licenses.capacity.measurement_unit<br><br>• Introduced in: 9.18 |
| licenses.capacity.maximum_size | integer | query | False | Filter by licenses.capacity.maximum_size |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| licenses.capacity.disabled_size | integer | query | False | Filter by licenses.capacity.disabled_size<br>• Introduced in: 9.18 |
| licenses.capacity.used_size | integer | query | False | Filter by licenses.capacity.used_size |
| licenses.serial_number | string | query | False | Filter by licenses.serial_number |
| licenses.evaluation | boolean | query | False | Filter by licenses.evaluation |
| licenses.shutdown_imminent | boolean | query | False | Filter by licenses.shutdown_imminent<br>• Introduced in: 9.11 |
| licenses.start_time | string | query | False | Filter by licenses.start_time |
| licenses.compliance.state | string | query | False | Filter by licenses.compliance.state |
| licenses.installed_license | string | query | False | Filter by licenses.installed_license<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| description | string | License description |
| entitlement | entitlement | |
| licenses | array[licenses] | Installed licenses of the package. |
| name | string | Name of the license. |
| scope | string | Scope of the license. |
| state | string | Summary state of package based on all installed licenses. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "NFS License",
  "entitlement": {
    "action": "string",
    "risk": "string"
  },
  "licenses": [
    {
      "capacity": {
        "disabled_size": 0,
        "maximum_size": 0,
        "measurement_unit": "string",
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 +0000",
      "host_id": "456-44-1234",
      "installed_license": "Core Bundle",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 +0000"
    }
  ],
  "name": "NFS",
  "scope": "string",
  "state": "compliant"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

entitlement

| Name | Type | Description |
|---|---|---|
| action | string | Entitlement action to be taken to mitigate the risk |
| risk | string | Entitlement risk of the package |

capacity

| Name | Type | Description |
|---|---|---|
| disabled_size | integer | Capacity that is currently disabled due to license limits. Its unit of measure is specified in the measurement_unit field. |
| maximum_size | integer | Licensed capacity size that can be used. Its unit of measure is specified in the measurement_unit field. |
| measurement_unit | string | Unit of measure for capacity based licenses. |
| used_size | integer | Specifies the total number of GPUs in the system when measurement_unit is GPUs, else specifies the bytes used. |

compliance

| Name | Type | Description |
|---|---|---|
| state | string | Compliance state of the license. |

licenses

| Name | Type | Description |
| --- | --- | --- |
| active | boolean | A flag indicating whether the license is currently being enforced. |
| capacity | capacity | |
| compliance | compliance | |
| evaluation | boolean | A flag indicating whether the license is in evaluation mode. |
| expiry_time | string | Date and time when the license expires. |
| host_id | string | A string that associates the license with a node or cluster. |
| installed_license | string | Name of license that enabled the feature. |
| owner | string | Cluster, node or license manager that owns the license. |
| serial_number | string | Serial number of the license. |
| shutdown_imminent | boolean | A flag indicating whether the Cloud ONTAP system is going to shutdown as the Cloud platform license has already expired.<br>• readOnly: 1<br>• Introduced in: 9.11<br>• x-nullable: true |
| start_time | string | Date and time when the license starts. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage the BlueXP cloud service

### Manage the BlueXP cloud service

#### Overview

You can use this API to ping the BlueXP cloud service. The POST operation retrieves the details about service reachability, configurability, and ping latency.

#### Performing a ping operation

You can perform a ping operation by issuing a POST request on /cluster/mediator-ping. Parameters are provided in the body of the POST request. There is only one required parameter.

### Required configuration fields

These fields are always required for any POST /cluster/mediator-ping request.

- `type` - Specifies the type of the mediator.

---

#### Examples

POST request body for a ping to the BlueXP cloud service.

```
# API
/api/cluster/mediator-ping
```

### POST request body from a file

```
ping_post_body.txt:
{
"type": "cloud"
}
curl -X POST https://<mgmt-ip>/api/cluster/mediator-ping -d
"@ping_post_body.txt"
```

#### Inline POST request body

```
curl -X POST https://<mgmt-ip>/api/cluster/mediator-ping -d
'{"type":"cloud"}'
```

#### POST request response

```
HTTP/1.1 200 OK
Cache-Control: no-cache,no-store,must-revalidate
Connection: close
Date: Mon, 03 Mar 2025 16:40:20 GMT
Server: libzapid-httpd
Vary: Accept-Encoding,Origin
Content-Length: 69
Content-Type: application/hal+json
Client-Date: Mon, 03 Mar 2025 16:40:20 GMT
Client-Peer: 10.235.144.44:80
Client-Response-Num: 1
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
X-Content-Type-Options: nosniff
{
"reachable": true,
"latency_ms": 50,
"configurable": true,
"high_latency": false,
"proxy_configured": true,
"proxy_used": true,
"timeout_occurred": false
}
```

#### Ping the BlueXP cloud service

POST `/cluster/mediator-ping`

**Introduced In:** 9.17

Pings BlueXP cloud service.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| configurable | boolean | Indicates if the BlueXP cloud mediator is configurable. This depends on whether the ping latency is within a threshold. |
| high_latency | boolean | Indicates if the ping latency of the BlueXP cloud server is greater than a threshold. |
| latency_ms | integer | Ping latency in milliseconds. |
| proxy_configured | boolean | Indicates if the HTTP proxy is configured on the cluster for the REST API calls to the BlueXP cloud server. |
| proxy_used | boolean | Indicates if the HTTP proxy is used for the ping to the BlueXP cloud server. |
| reachable | boolean | Ping status of the BlueXP cloud service. |
| timeout_occurred | boolean | Indicates if the ping to the BlueXP cloud server failed due to a timeout. |
| type | string | Mediator type. |

#### Example request

```
{
  "configurable": 1,
  "high_latency": "",
  "latency_ms": 150,
  "proxy_configured": 1,
  "proxy_used": 1,
  "reachable": 1,
  "timeout_occurred": "",
  "type": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| configurable | boolean | Indicates if the BlueXP cloud mediator is configurable. This depends on whether the ping latency is within a threshold. |
| high_latency | boolean | Indicates if the ping latency of the BlueXP cloud server is greater than a threshold. |
| latency_ms | integer | Ping latency in milliseconds. |
| proxy_configured | boolean | Indicates if the HTTP proxy is configured on the cluster for the REST API calls to the BlueXP cloud server. |
| proxy_used | boolean | Indicates if the HTTP proxy is used for the ping to the BlueXP cloud server. |
| reachable | boolean | Ping status of the BlueXP cloud service. |
| timeout_occurred | boolean | Indicates if the ping to the BlueXP cloud server failed due to a timeout. |

| Name | Type | Description |
|---|---|---|
| type | string | Mediator type. |

**Example response**

```
{
  "configurable": 1,
  "high_latency": "",
  "latency_ms": 150,
  "proxy_configured": 1,
  "proxy_used": 1,
  "reachable": 1,
  "timeout_occurred": "",
  "type": "string"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 13369414 | Failed to talk to the BlueXP cloud service. Check the need for an HTTP proxy and network settings like firewall. |
| 13369417 | Failed to check if an HTTP proxy is configured on the cluster. |

| Name | Type | Description |
|---|---|---|
| configurable | boolean | Indicates if the BlueXP cloud mediator is configurable. This depends on whether the ping latency is within a threshold. |
| error | error | |
| reachable | boolean | Ping status of the BlueXP cloud service. |

### Example error

```
{
  "configurable": "",
  "error": {
    "message": "string"
  },
  "reachable": ""
}
```

### Definitions

**See Definitions**

mediator_ping

BlueXP cloud service ping API.

| Name | Type | Description |
|---|---|---|
| configurable | boolean | Indicates if the BlueXP cloud mediator is configurable. This depends on whether the ping latency is within a threshold. |
| high_latency | boolean | Indicates if the ping latency of the BlueXP cloud server is greater than a threshold. |
| latency_ms | integer | Ping latency in milliseconds. |
| proxy_configured | boolean | Indicates if the HTTP proxy is configured on the cluster for the REST API calls to the BlueXP cloud server. |
| proxy_used | boolean | Indicates if the HTTP proxy is used for the ping to the BlueXP cloud server. |
| reachable | boolean | Ping status of the BlueXP cloud service. |
| timeout_occurred | boolean | Indicates if the ping to the BlueXP cloud server failed due to a timeout. |
| type | string | Mediator type. |

error

| Name | Type | Description |
|---|---|---|
| code | integer | Error code. |
| message | string | Error message string. |

## Manage cluster mediators

## Manage cluster mediators

### Overview

You can use this API to add, modify or remove a mediator in a MetroCluster IP configuration or a SnapMirror active sync configuration. You can also use this API to get the status and details of the existing mediator. The GET operation returns the status of the mediator along with the mediator details. The DELETE operation removes the mediator. The POST operation adds the mediator. The PATCH operation modifies the local and remote proxy options.
SnapMirror active sync supports two types of mediators:

1. **ONTAP Mediator:** This is the traditional mediator used in SnapMirror active sync, and is hosted on-premises. It requires deployment at a third, neutral site.
2. **ONTAP cloud mediator:** This mediator is hosted in BlueXP SaaS and is primarily designed for SnapMirror active sync. It eliminates the need for a third site and provides enhanced scalability and availability.

### Adding a mediator

A mediator can be added by issuing a POST request on /cluster/mediators. Parameters are provided in the body of the POST request. There are no optional parameters for adding a mediator to a MetroCluster IP configuration.

#### Required configuration fields (MetroCluster)

These fields are always required for any POST /cluster/mediators request.

- `ip_address` - Specifies the IP address of the mediator.
- `user` - Specifies a user name credential.
- `password` - Specifies a password credential.

#### Required configuration fields (SnapMirror active sync: ONTAP Mediator, ONTAP cloud mediator)

These fields are required for any POST /cluster/mediators request.

- `type` - (defaults to "on-prem" if omitted) Specifies the type of mediator. For the ONTAP Mediator, the value of type has to be "on-prem". For the ONTAP cloud mediator, the value of type has to be "cloud".
- `peer_cluster.name` - Specifies the name of the peer cluster.
- `ip_address` - (only applicable to the ONTAP Mediator) Specifies the IP address of the mediator.
- `user` - (only applicable to the ONTAP Mediator) Specifies the user name credential.
- `password` - (only applicable to the ONTAP Mediator) Specifies a password credential.
- `ca_certificate` - (optional if the certificate is already installed, only applicable to the ONTAP Mediator) Specifies the CA certificate for the ONTAP Mediator.
- `bluexp_org_id` - (only applicable to the ONTAP cloud mediator) Specifies the BlueXP organization ID.
- `service_account_client_id` - (only applicable to the ONTAP cloud mediator) Specifies the client ID of the service account.
- `service_account_client_secret` - (only applicable to the ONTAP cloud mediator) Specifies the client secret of the service account.

- `bluexp_account_token` - (only applicable to the ONTAP cloud mediator) Specifies the BlueXP service account token. This field is mutually exclusive with the `service_account_client_id` and `service_account_client_secret` pair, meaning either the token or the client-id and client-secret pair is allowed.
- `use_http_proxy_local` - (optional, defaults to false if omitted, only applicable for ONTAP cloud mediator) Specifies if a HTTP proxy should be used on the ONTAP cluster.
- `use_http_proxy_remote` - (optional, defaults to false if omitted, only applicable for ONTAP cloud mediator) Specifies if a HTTP proxy should be used on the peer ONTAP cluster.
- `strict_cert_validation` - (optional, defaults to false if omitted, only applicable for ONTAP cloud mediator) Specifies if strict validation of certificates is performed while making REST API calls to the ONTAP Cloud Mediator.

### Polling the setup job

After a successful POST /cluster/mediators is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The setup job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

#### Deleting a mediator

A mediator can be deleted by issuing a DELETE to /cluster/mediators/{uuid}. Parameters are provided in the body of the DELETE request. There are no optional parameters for deleting a mediator in a MetroCluster IP configuration.

### Required configuration fields (MetroCluster)

These fields are always required for any DELETE /cluster/mediators/{uuid} request.

- `user` - Specifies a user name credential.
- `password` - Specifies a password credential.

### Required configuration fields (SnapMirror active sync: ONTAP Mediator, ONTAP cloud mediator)

No fields are required in the body of the DELETE /cluster/mediators/{uuid} request for a mediator in SnapMirror active sync configuration.

### Polling the delete job

After a successful DELETE /cluster/mediators/{uuid} is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The delete job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

#### Modify a mediator

An ONTAP cloud mediator can be modified by issuing a PATCH to /cluster/mediators/{uuid}. Parameters are provided in the body of the PATCH request. Currently there are only 2 parameters which can be modified.

### Fields which can be modified

These fields can be modified by any PATCH /cluster/mediators/{uuid} request.

- `use-http-proxy-local` - Specifies if the local cluster should use http-proxy server while making REST API calls to mediator.
- `use-http-proxy-remote` - Specifies if the remote cluster should use http-proxy while making REST API calls to mediator.

### Polling the modify job

After a successful PATCH /cluster/mediators/{uuid} is issued, an HTTP status code of 202 (Accepted) is returned along with a link in the body of the response. The modify response can be monitored by using the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

---

#### Examples

### Setting up a mediator for a four-node MetroCluster and a SnapMirror active sync configuration

This example shows the POST body when setting up a mediator for a four-node MetroCluster IP configuration and a SnapMirror active sync configuration. For MetroCluster, the only prerequisite is that MetroCluster IP is configured.

```
# API
/api/cluster/mediators
```

### POST body included from file (MetroCluster)

```
mediator_post_body.txt:
{
"ip_address": "1.1.1.1",
"user": "username",
"password": "password"
}
curl -X POST https://<mgmt-ip>/api/cluster/mediators -d
"@mediator_post_body.txt"
```

### POST body included from file (SnapMirror active sync: ONTAP Mediator)

```
mediator_post_body.txt:
{
"peer_cluster.name": "C2_sti230-vsim-sr092w_cluster",
"ip_address": "172.18.48.61",
"user": "username",
"password": "password"
}
curl -X POST https://<mgmt-ip>/api/cluster/mediators -d
"@mediator_post_body.txt"
```

**POST body included from file (SnapMirror active sync: ONTAP cloud mediator)**

```
mediator_post_body.txt:
{
"peer_cluster.name": "C2_sti230-vsim-sr092w_cluster",
"type": "cloud",
"bluexp_org_id":"your-bluexp-org-id",
"service_account_client_id":"your-account-client-id",
"service_account_client_secret":"your-account-client-secret",
"use_http_proxy_local":"true",
"use_http_proxy_remote":"true",
"strict_cert_validation":"true"
}
curl -X POST https://<mgmt-ip>/api/cluster/mediators -d
"@mediator_post_body.txt"
```

**Inline POST body (MetroCluster)**

```
curl -X POST https://<mgmt-ip>/api/cluster/mediators -H "Content-Type:
application/hal+json" -d '{"ip_address": "1.1.1.1", "user": "username",
"password": "password"}'
```

**Inline POST body (SnapMirror active sync: ONTAP Mediator)**

```
curl -X POST https://<mgmt-ip>/api/cluster/mediators -H "Content-Type:
application/hal+json" -d '{"peer_cluster.name": "C2_sti230-vsim-
sr092w_cluster", "ip_address": "172.18.48.61", "user": "username",
"password": "password"}'
```

**Inline POST body (SnapMirror active sync: ONTAP cloud mediator)**

```
curl -X POST https://<mgmt-ip>/api/cluster/mediators -H "Content-Type:
application/hal+json" -d '{"peer_cluster.name": "C2_sti230-vsim-
sr092w_cluster", "type": "cloud", "bluexp_org_id": "your-bluexp-org-id",
"service_account_client_id": "your-account-client-id",
"service_account_client_secret": "your-account-client-secret",
"use_http_proxy_local": "true", "use_http_proxy_remote": "true",
"strict_cert_validation": "true"}'
```

**POST Response**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 07:40:59 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "f567b48b-fca6-11ea-acaf-005056bb47c1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f567b48b-fca6-11ea-acaf-005056bb47c1"
    }
  }
}
}
```

**Monitoring the job progress**

Use the link provided in the response to the POST request to fetch information for the mediator setup job.

**Request**

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/f567b48b-fca6-11ea-acaf-
005056bb47c1
```

**Job status response**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 07:41:29 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
"uuid": "f567b48b-fca6-11ea-acaf-005056bb47c1",
"description": "POST /api/cluster/mediators/",
"state": "running",
"start_time": "2020-09-22T03:41:00-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f567b48b-fca6-11ea-acaf-005056bb47c1"
  }
}
}
```

**Final status of a successful Mediator add**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 07:43:38 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 358
Content-Type: application/hal+json
{
"uuid": "f567b48b-fca6-11ea-acaf-005056bb47c1",
"description": "POST /api/cluster/mediators/",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-09-22T03:41:00-04:00",
"end_time": "2020-09-22T03:42:10-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f567b48b-fca6-11ea-acaf-005056bb47c1"
  }
}
}
```

#### Retrieving the existing mediator configurations

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/mediators
```

#### Response

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 08:53:18 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 320
Content-Type: application/hal+json
{
"records": [
  {
    "uuid": "f89e8906-fca6-11ea-acaf-005056bb47c1",
    "_links": {
      "self": {
        "href": "/api/cluster/mediators/f89e8906-fca6-11ea-acaf-
005056bb47c1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/cluster/mediators"
  }
}
}
```

#### Retrieving a specific mediator using the UUID

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/mediators/f89e8906-fca6-11ea-
acaf-005056bb47c1
```

**Response (MetroCluster)**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 08:59:40 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 347
Content-Type: application/hal+json
{
"uuid": "f89e8906-fca6-11ea-acaf-005056bb47c1",
"ip_address": "10.234.173.40",
"port": 31784,
"reachable": true,
"peer_cluster": {
  "name": "mcc_siteB",
  "uuid": "38779fd1-fc6b-11ea-9421-005056bb21d8"
},
"peer_mediator_connectivity": "connected",
"strict_cert_validation": "true",
"_links": {
  "self": {
    "href": "/api/cluster/mediators/f89e8906-fca6-11ea-acaf-005056bb47c1"
  }
}
}
```

**Response (SnapMirror active sync: ONTAP Mediator)**

```
{
"uuid": "cc44f61f-ffd5-11ef-aaa6-005056ae32ff",
"ip_address": "172.18.48.61",
"port": 31784,
"reachable": true,
"peer_cluster": {
  "name": "C2_sti232-vsim-sr089o_cluster",
  "uuid": "ece8d7c6-fd8f-11ef-b8b3-005056aec21e"
},
"peer_mediator_connectivity": "connected",
"local_mediator_connectivity": "connected",
"type": "on_prem",
"_links": {
  "self": {
    "href": "/api/cluster/mediators/cc44f61f-ffd5-11ef-aaa6-005056ae32ff"
  }
}
}
```


**Response (SnapMirror active sync: ONTAP cloud mediator)**


```
{
"uuid": "b9ba4e5b-ff14-11ef-82a4-005056ae1bc8",
"ip_address": "0.0.0.0",
"reachable": true,
"peer_cluster": {
  "name": "C2_sti232-vsim-sr089o_cluster",
  "uuid": "ece8d7c6-fd8f-11ef-b8b3-005056aec21e"
},
"peer_mediator_connectivity": "connected",
"local_mediator_connectivity": "connected",
"use_http_proxy_local": true,
"use_http_proxy_remote": true,
"type": "cloud",
"_links": {
  "self": {
    "href": "/api/cluster/mediators/b9ba4e5b-ff14-11ef-82a4-005056ae1bc8"
  }
}
}
```


**Deleting a configured mediator using the UUID (MetroCluster)**

### Request

```
curl -X DELETE https://<mgmt-ip>/api/cluster/mediators/{uuid} -H "Content-
Type: application+hal/json" -d '{"user": "username", "password":
"password"}'
```

### Response

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 09:13:52 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "eeb71ccd-fcb3-11ea-acaf-005056bb47c1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/eeb71ccd-fcb3-11ea-acaf-005056bb47c1"
    }
  }
}
}
```

#### Deleting a configured mediator using the UUID (SnapMirror active sync: ONTAP Mediator, ONTAP cloud mediator)

### Request

```
curl -X DELETE https://<mgmt-ip>/api/cluster/mediators/{uuid} -H "Content-
Type: application+hal/json"
```

### Response

```
{
"job": {
  "uuid": "0661a77a-ff16-11ef-82a4-005056ae1bc8",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/0661a77a-ff16-11ef-82a4-005056ae1bc8"
    }
  }
}
}
```

#### Monitoring the job progress

Use the link provided in the response to the DELETE request to fetch information for the delete job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/eeb71ccd-fcb3-11ea-acaf-
005056bb47c1
```

#### Job status response

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 09:14:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 316
Content-Type: application/hal+json
{
"uuid": "eeb71ccd-fcb3-11ea-acaf-005056bb47c1",
"description": "DELETE /api/cluster/mediators/f89e8906-fca6-11ea-acaf-
005056bb47c1",
"state": "running",
"start_time": "2020-09-22T05:13:52-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/eeb71ccd-fcb3-11ea-acaf-005056bb47c1"
  }
}
}
```

#### Final status of the Mediator DELETE job

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 09:21:46 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 396
Content-Type: application/hal+json
{
"uuid": "eeb71ccd-fcb3-11ea-acaf-005056bb47c1",
"description": "DELETE /api/cluster/mediators/f89e8906-fca6-11ea-acaf-
005056bb47c1",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-09-22T05:13:52-04:00",
"end_time": "2020-09-22T05:14:24-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/eeb71ccd-fcb3-11ea-acaf-005056bb47c1"
  }
}
}
```

#### Modifying a configured Mediator using the uuid

#### Request

```
curl -X PATCH https://<mgmt-ip>/api/cluster/mediators/{uuid} -H "Content-
Type: application+hal/json" -d '{"use_http_proxy_local":"true",
"use_http_proxy_remote":"true", "strict_cert_validation": "true"}'
```

#### Response

```
HTTP/1.1 202 Accepted
Date: Thu, 06 Mar 2025 10:10:10 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "3016e8e9-fa73-11ef-9d41-005056ae10ae",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3016e8e9-fa73-11ef-9d41-005056ae10ae"
    }
  }
}
}
```

#### Monitoring the job progress

Use the link provided in the response to the PATCH request to fetch information for the modify job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/3016e8e9-fa73-11ef-9d41-
005056ae10ae
```

#### Job status response

```
HTTP/1.1 202 Accepted
Date: Thu, 06 Mar 2025 10:10:22 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 395
Content-Type: application/hal+json
{
"uuid": "3016e8e9-fa73-11ef-9d41-005056ae10ae",
"description": "PATCH /api/cluster/mediators/a9a812bf-f7f8-11ef-a923-
005056ae10ae",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2025-03-06T05:10:10-05:00",
"end_time": "2025-03-06T05:10:10-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/3016e8e9-fa73-11ef-9d41-005056ae10ae"
  }
}
}
```

### Retrieve ONTAP Mediators configured in the cluster

GET `/cluster/mediators`

**Introduced In:** 9.8

"Retrieves a Mediator configured in the cluster."

#### Related ONTAP commands

- `storage iscsi-initiator show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| peer_mediator_connectivity | string | query | False | Filter by peer_mediator_connectivity<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| use_http_proxy_local | boolean | query | False | Filter by use_http_proxy_local<br>• Introduced in: 9.17 |
| type | string | query | False | Filter by type<br>• Introduced in: 9.17 |
| port | integer | query | False | Filter by port |
| uuid | string | query | False | Filter by uuid |
| strict_cert_validation | boolean | query | False | Filter by strict_cert_validation<br>• Introduced in: 9.17 |
| peer_cluster.name | string | query | False | Filter by peer_cluster.name |
| peer_cluster.uuid | string | query | False | Filter by peer_cluster.uuid |
| ip_address | string | query | False | Filter by ip_address |
| local_mediator_connectivity | string | query | False | Filter by local_mediator_connectivity<br>• Introduced in: 9.17 |
| reachable | boolean | query | False | Filter by reachable |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "dr_group": {
        "id": 0
      },
      "ip_address": "10.10.10.7",
      "local_mediator_connectivity": "connected",
      "peer_cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster2",
        "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
      },
      "peer_mediator_connectivity": "connected",
      "port": 31784,
      "reachable": 1,
      "strict_cert_validation": 1,
      "type": "string",
      "use_http_proxy_local": 1,
      "use_http_proxy_remote": 1,
      "uuid": "string"
    }
  ]
}
```


### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 2430739 | Unable to access Mediator. Reason: Invalid Mediator IP or Mediator does not exist. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

records

Mediator information

| Name | Type | Description |
|---|---|---|
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| uuid | string | The unique identifier for the mediator service. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create and connect an ONTAP Mediator**

POST `/cluster/mediators`

**Introduced In:** 9.8

Creates and connect a mediator.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| bluexp_account_token | string | BlueXP account token. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| bluexp_org_id | string | BlueXP organization ID. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| ca_certificate | string | CA certificate for ONTAP Mediator. This is optional if the certificate is already installed.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.8<br>• x-nullable: true |
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| password | string | The password used to connect to the REST server on the mediator. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |

| Name | Type | Description |
| --- | --- | --- |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| service_account_client_id | string | Client ID of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| service_account_client_secret | string | Client secret token of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| user | string | The username used to connect to the REST server on the mediator. |
| uuid | string | The unique identifier for the mediator service. |

#### Example request

```
{
  "bluexp_account_token": "string",
  "bluexp_org_id": "string",
  "ca_certificate": "string",
  "dr_group": {
    "id": 0
  },
  "ip_address": "10.10.10.7",
  "local_mediator_connectivity": "connected",
  "password": "mypassword",
  "peer_cluster": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "peer_mediator_connectivity": "connected",
  "port": 31784,
  "reachable": 1,
  "service_account_client_id": "string",
  "service_account_client_secret": "string",
  "strict_cert_validation": 1,
  "type": "string",
  "use_http_proxy_local": 1,
  "use_http_proxy_remote": 1,
  "user": "myusername",
  "uuid": "string"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 13369351 | Update to mediator failed. Reason: does not authorized for that command. Check that the peer cluster and mediator are reachable. |
| 2430758 | Mediator cannot be added. Reason: Some of the nodes in the cluster have incorrect MetroCluster IP configuration. |
| 2430734 | Unable to add Mediator. Reason : Authentication failed. |
| 2430774 | Failed to add Mediator. Reason : Node is not reachable or does not exist. |
| 2432871 | Failed to add Mediator. Reason : Some of the parameters in the request is invalid. |

#### Definitions

**See Definitions**

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
|---|---|---|
| bluexp_account_token | string | BlueXP account token. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| bluexp_org_id | string | BlueXP organization ID. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| ca_certificate | string | CA certificate for ONTAP Mediator. This is optional if the certificate is already installed.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.8<br>• x-nullable: true |
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| password | string | The password used to connect to the REST server on the mediator. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| service_account_client_id | string | Client ID of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| service_account_client_secret | string | Client secret token of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| user | string | The username used to connect to the REST server on the mediator. |
| uuid | string | The unique identifier for the mediator service. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an ONTAP Mediator**

DELETE `/cluster/mediators/{uuid}`

**Introduced In:** 9.8

Deletes the mediator.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | • format: uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| bluexp_account_token | string | BlueXP account token. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| bluexp_org_id | string | BlueXP organization ID. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| ca_certificate | string | CA certificate for ONTAP Mediator. This is optional if the certificate is already installed.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.8<br>• x-nullable: true |
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| password | string | The password used to connect to the REST server on the mediator. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| service_account_client_id | string | Client ID of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| service_account_client_secret | string | Client secret token of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| user | string | The username used to connect to the REST server on the mediator. |
| uuid | string | The unique identifier for the mediator service. |

#### Example request

```
{
  "bluexp_account_token": "string",
  "bluexp_org_id": "string",
  "ca_certificate": "string",
  "dr_group": {
    "id": 0
  },
  "ip_address": "10.10.10.7",
  "local_mediator_connectivity": "connected",
  "password": "mypassword",
  "peer_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "peer_mediator_connectivity": "connected",
  "port": 31784,
  "reachable": 1,
  "service_account_client_id": "string",
  "service_account_client_secret": "string",
  "strict_cert_validation": 1,
  "type": "string",
  "use_http_proxy_local": 1,
  "use_http_proxy_remote": 1,
  "user": "myusername",
  "uuid": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 13369377 | Mediator field "mediator.id" does not exist. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
|---|---|---|
| bluexp_account_token | string | BlueXP account token. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| bluexp_org_id | string | BlueXP organization ID. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| ca_certificate | string | CA certificate for ONTAP Mediator. This is optional if the certificate is already installed.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.8<br>• x-nullable: true |
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| password | string | The password used to connect to the REST server on the mediator. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| service_account_client_id | string | Client ID of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| service_account_client_secret | string | Client secret token of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| user | string | The username used to connect to the REST server on the mediator. |
| uuid | string | The unique identifier for the mediator service. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the ONTAP Mediator state and configuration

GET `/cluster/mediators/{uuid}`

**Introduced In:** 9.8

"Retrieves the Mediator state and configuration."

#### Related ONTAP commands

- `storage iscsi-initiator show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | • format: uuid |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| dr_group | dr_group | DR group reference. |

| Name | Type | Description |
| --- | --- | --- |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| uuid | string | The unique identifier for the mediator service. |

**Example response**


```
{
  "dr_group": {
    "id": 0
  },
  "ip_address": "10.10.10.7",
  "local_mediator_connectivity": "connected",
  "peer_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "peer_mediator_connectivity": "connected",
  "port": 31784,
  "reachable": 1,
  "strict_cert_validation": 1,
  "type": "string",
  "use_http_proxy_local": 1,
  "use_http_proxy_remote": 1,
  "uuid": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Modify the mediator configuration

PATCH `/cluster/mediators/{uuid}`

**Introduced In:** 9.17

Modifies the mediator configuration.

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | • format: uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| dr_group | dr_group | DR group reference. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |

| Name | Type | Description |
| --- | --- | --- |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier for the mediator service. |

#### Example request

```
{
  "dr_group": {
    "id": 0
  },
  "local_mediator_connectivity": "connected",
  "peer_mediator_connectivity": "connected",
  "reachable": 1,
  "strict_cert_validation": 1,
  "use_http_proxy_local": 1,
  "use_http_proxy_remote": 1,
  "uuid": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 13369377 | Mediator field "mediator.id" does not exist. |
| 13369412 | Modification of the mediator configuration succeeded on the peer cluster but failed on the local cluster. Reason : {Reason of failure}. Wait a few minutes, then try the operation again. |
| 13369421 | Failed to modify mediator configuration on the cluster. Reason : {Reason of failure}. Wait a few minutes, then try the operation again. |

#### Definitions

**See Definitions**

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
|---|---|---|
| dr_group | dr_group | DR group reference. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |

| Name | Type | Description |
| --- | --- | --- |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| uuid | string | The unique identifier for the mediator service. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for the cluster**

GET `/cluster/metrics`

**Introduced In:** 9.6

Retrieves historical performance metrics for the cluster.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| throughput.other | integer | query | False | Filter by throughput.other |
| throughput.total | integer | query | False | Filter by throughput.total |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| timestamp | string | query | False | Filter by timestamp |
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, or 1y.<br>The period for each time range is specified as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and manage MetroCluster configurations

### View and manage MetroCluster configurations

### Overview

You can use this API to create, perform operations, and retrieve relevant information pertaining to MetroCluster. The GET operation fetches MetroCluster status and configuration parameters for the local and partner cluster. The PATCH operation executes a switchover, heal or switchback operation. The POST request can be used to setup a MetroCluster.

#### Creating a MetroCluster

A new MetroCluster can be set up by issuing a POST to /cluster/metrocluster. Parameters are provided in the body of the POST request.

### Fields used for setting up a MetroCluster configuration

The fields used for MetroCluster APIs are either required or optional and are described as follows:

### Required configuration fields

These fields are always required for any POST /cluster/metrocluster request.

- `partner_cluster.name` - Specifies the partner cluster name to which cluster peering has been established.
- `dr_pairs` - Specifies local and DR partner node pairs. Each pair uniquely identifies a DR group.

### Optional configuration fields

This field is used to set up additional components in a MetroCluster configuration.

- `mediator.*` - Specifies mediator parameters. If Mediator Assisted Unplanned Switchover (MAUSO) functionality is required, then a mediator should be configured.
- `mccip_ports` - Specifies relevant layer 3 network configuration information for each port. These include port name, node name, IP address, gateway, and netmask. If mccip_ports is not provided, then the API automatically generates IP addresses for the ports and creates a layer 2 network configuration.

### Polling the setup job

After a successful POST /cluster/metrocluster is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The setup job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

---

#### Examples

##### Setting up a 4-node MetroCluster

```
This example shows the POST body when setting up a 4-node MetroCluster
along with a mediator. It is required that cluster peering be established
between two clusters, in this example, site "mcc_siteA" and "mcc_siteB"
before issuing the POST request. Nodes "node-a" and "node-b" are HA
partners and part of the local cluster "mcc_siteA", wheres nodes "node-c"
and "node-d" are HA partners in the partner cluster  "mcc_siteB".
Specifying a single DR pairing of "node-a" and "node-c" is sufficient to
identify a DR group -- "node-a" and "node-c" will be designated primary DR
partners ("node-b" and "node-d" too). "node-d" will then be designated
auxiliary partner of "node-a". Once the MetroCluster configuration has
been completed, and since mediator parameters have been provided, the
mediator will be setup and MAUSO enabled.

# API
/api/cluster/metrocluster
```

### POST body included from file

```
mcc_post_body.txt:
{
"partner_cluster" : {
  "name": "mcc_siteB"
},
"dr_pairs" : [
  {
    "node" : {
      "name" : "node-a"
    },
    "partner" : {
      "name" : "node-c"
    }
  }
],
"mediator" :  {
  "ip_address" : "1.2.3.4",
  "user" : "mcc_mediator",
  "password" : "openMediator"
}
}
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster -d
"@mcc_post_body.txt"
```

**Inline POST body**

```
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster -d
'{"partner_cluster" : {"name": "mcc_siteB" }, "dr_pairs" : [{"node" :
{"name" : "node-a" }, "partner" : {"name" : "node-c" }}], "mediator" :
{"ip_address" : "1.2.3.4", "user" : "mcc_mediator" , "password" :
"openMediator" }}'
```

**POST Response**

```
HTTP/1.1 202 Accepted
Date: Thu, 09 Jan 2020 20:38:05 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "f23abbdb-331f-11ea-acd3-005056a708b2",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
    }
  }
}
}
```

#### Monitoring the job progress

Use the link provided in the response to the POST request to fetch information for the setup job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/f23abbdb-331f-11ea-acd3-
005056a708b2
```

#### Job status response

The following is an example of the job status response returned by the running MetroCluster setup job:

```
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:40:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 373
Content-Type: application/hal+json
{
"uuid": "f23abbdb-331f-11ea-acd3-005056a708b2",
"description": "POST /api/cluster/metrocluster",
"state": "running",
"message": "Checking remote storage pool",
"code": 2432844,
"start_time": "2020-01-09T15:38:08-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
  }
}
```

#### Completion message

This is the final update message from the setup job indicating completion.

```
{
"uuid": "f23abbdb-331f-11ea-acd3-005056a708b2",
"description": "POST /api/cluster/metrocluster",
"state": "running",
"message": "MetroCluster setup is complete",
"code": 2432849,
"start_time": "2020-01-09T15:38:08-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
  }
}
}
```

#### Final status of a successful MetroCluster setup workflow

When the setup job completes, the 'end_time' field is populated, and the 'state' and 'message' fields report the final status.

```
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:43:54 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 360
Content-Type: application/hal+json
{
"uuid": "f23abbdb-331f-11ea-acd3-005056a708b2",
"description": "POST /api/cluster/metrocluster",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-01-09T15:38:08-05:00",
"end_time": "2020-01-09T15:43:50-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
  }
}
}
```

**Retrieving the MetroCluster configuration after completion of the POST request**

**Request**

```
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster
```

**Response**

```
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:49:40 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 849
Content-Type: application/hal+json
{
"local": {
  "configuration_state": "configured",
  "periodic_check_enabled": true,
  "mode": "normal",
  "partner_cluster_reachable": true,
  "cluster": {
    "name": "mcc_siteA",
    "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
    "_links": {
      "self": {
        "href": "/api/cluster"
      }
    }
  }
},
"remote": {
  "configuration_state": "configured",
  "periodic_check_enabled": true,
  "mode": "normal",
  "cluster": {
    "name": "mcc_siteB",
    "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/4207c6a5-30e2-11ea-be25-
005056a7dc84/cluster"
      }
    }
  }
},
"configuration_type": "ip_fabric",
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster"
  }
}
}
```

#### Retrieving information about the nodes in a MetroCluster configuration

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster/nodes
```

#### Response

```
HTTP/1.1 200 OK
Date: Fri, 10 Jan 2020 02:26:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Transfer-Encoding: chunked
{
"records": [
  {
    "dr_group_id": 1,
    "cluster": {
      "name": "mcc_siteA",
      "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
      "_links": {
        "self": {
          "href": "/api/cluster"
        }
      }
    },
    "node": {
      "name": "node-a",
      "uuid": "1e6b0137-30dd-11ea-82ba-005056a7c78a",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/1e6b0137-30dd-11ea-82ba-
005056a7c78a"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/nodes/1e6b0137-30dd-11ea-82ba-
005056a7c78a"
      }
    }
  },
```

```
  {
    "dr_group_id": 1,
    "cluster": {
      "name": "mcc_siteA",
      "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
      "_links": {
        "self": {
          "href": "/api/cluster"
        }
      }
    },
    "node": {
      "name": "node-b",
      "uuid": "1e57ba22-30dd-11ea-8b19-005056a708b2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/1e57ba22-30dd-11ea-8b19-
005056a708b2"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/nodes/1e57ba22-30dd-11ea-8b19-
005056a708b2"
      }
    }
  },
  {
    "dr_group_id": 1,
    "cluster": {
      "name": "mcc_siteB",
      "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
      "_links": {
        "self": {
          "href": "/api/cluster/peers/4207c6a5-30e2-11ea-be25-
005056a7dc84/cluster"
        }
      }
    },
    "node": {
      "name": "node-c",
      "uuid": "1e563efc-30dd-11ea-a9d3-005056a71573",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/1e563efc-30dd-11ea-a9d3-
```

```
005056a71573"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/nodes/1e563efc-30dd-11ea-a9d3-
005056a71573"
      }
    }
  },
  {
    "dr_group_id": 1,
    "cluster": {
      "name": "mcc_siteB",
      "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
      "_links": {
        "self": {
          "href": "/api/cluster/peers/4207c6a5-30e2-11ea-be25-
005056a7dc84/cluster"
        }
      }
    },
    "node": {
      "name": "node-d",
      "uuid": "1e400aa4-30dd-11ea-adec-005056a7dc84",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/1e400aa4-30dd-11ea-adec-
005056a7dc84"
        }
      }
    },
    "encryption_enabled": false,
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/nodes/1e400aa4-30dd-11ea-adec-
005056a7dc84"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/nodes"
```

```
  }
}
}
```

### Retrieving MetroCluster status and configuration information

```
GET https://<mgmt-ip>/api/cluster/metrocluster
{
  "local": {
      "configuration_state": "configured",
      "periodic_check_enabled": true,
      "mode": "normal",
      "cluster": {
          "name": "cluster1",
          "uuid": "bbc00ca3-8d81-11e9-b5a9-005056826931",
          "_links": {
              "self": {
                  "href": "/api/cluster"
              }
          }
      }
  },
  "remote": {
      "configuration_state": "configured",
      "periodic_check_enabled": true,
      "mode": "normal",
      "cluster": {
          "name": "cluster3",
          "uuid": "ce2cf803-8d81-11e9-87db-00505682cecf",
          "_links": {
              "self": {
                  "href": "/api/cluster/peers/ce2cf803-8d81-11e9-87db-
00505682cecf/cluster"
              }
          }
      }
  },
  "_links": {
      "self": {
          "href": "/api/cluster/metrocluster"
      }
  }
}
```

### Initiating a switchover, heal or switchback command using PATCH

PATCH is used to initiate a variety of operations by specifying one of the following values in the "action" parameter:

- `switchover` - Initiates an Unplanned Switchover (USO).
- `negotiated_switchover` - Indicates that an Negotiated switchover (NSO) is to be performed.
- `negotiated_switchover_simulate` - Provides validation in preparation for NSO but does not perform the operation.
- `switchback` - Indicates that a switchback is to be performed.
- `switchback_simulate` - Provides validation for switchback but does not commit the operation.
- `heal_aggregates` - Indicates that the aggregates phase of the heal operation is to be performed.
- `heal_root_aggregates` - Indicates that the root aggregates phase of the heal operation is to be performed.

#### PATCH Switchover example

```
PATCH https://<mgmt-ip>/api/cluster/metrocluster?action=switchover
{
  "job": {
      "uuid": "70e54274-57ee-11e9-aa33-005056820b99",
      "_links": {
          "self": {
              "href": "/api/cluster/jobs/70e54274-57ee-11e9-aa33-
005056820b99"
          }
      }
  }
}
```

This returns a job UUID. A subsequent GET for this job should return the following:

```
GET https://<mgmt-ip>/api/cluster/jobs/70e54274-57ee-11e9-aa33-
005056820b99
{
  "uuid": "70e54274-57ee-11e9-aa33-005056820b99",
  "description": "MetroCluster Switchover Job",
  "state": "success",
  "message": "Complete: Switchover is successful.",
  "code": 0,
  "start_time": "2019-04-05T15:02:02-07:00",
  "end_time": "2019-04-05T15:02:30-07:00",
  "_links": {
      "self": {
          "href": "/api/cluster/jobs/70e54274-57ee-11e9-aa33-005056820b99"
      }
  }
}
```

**PATCH Switchback example:**

```
PATCH https://<mgmt-ip>/api/cluster/metrocluster?action=switchback
{
  "job": {
      "uuid": "a62714cc-57ec-11e9-aa33-005056820b99",
      "_links": {
          "self": {
              "href": "/api/cluster/jobs/a62714cc-57ec-11e9-aa33-
005056820b99"
          }
      }
  }
}
```

This returns a job UUID with a link to the job. A subsequent GET for this job UUID can be used to retrieve the completion status of the operation:

```
GET https://<mgmt-ip>/api/cluster/jobs/a62714cc-57ec-11e9-aa33-
005056820b99
{
  "uuid": "a62714cc-57ec-11e9-aa33-005056820b99",
  "description": "MetroCluster Switchback Job",
  "state": "success",
  "message": "Complete: Switchback is successful.",
  "code": 0,
  "start_time": "2019-04-05T14:49:12-07:00",
  "end_time": "2019-04-05T14:50:12-07:00",
  "_links": {
      "self": {
          "href": "/api/cluster/jobs/a62714cc-57ec-11e9-aa33-005056820b99"
      }
  }
}
```

### Retrieve MetroCluster status and configuration details

GET `/cluster/metrocluster`

**Introduced In:** 9.8

Retrieves MetroCluster status and configuration details.

#### Related ONTAP commands

- `metrocluster show`
- `metrocluster node show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| configuration_type | string | Displays the MetroCluster configuration type. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| local | local | |
| remote | remote | |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configuration_type": "string",
  "local": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  },
  "remote": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_pairs

| Name | Type | Description |
|---|---|---|
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

cluster

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

local

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the local cluster configuration. |
| mode | string | Specifies the mode of operation of the local cluster. |
| partner_cluster_reachable | boolean | Specifies whether the partner cluster is reachable from the local cluster. |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the local cluster. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
|---|---|---|
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
|---|---|---|
| l3_config | l3_config | |
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

dr_group

DR group reference.

| Name | Type | Description |
|---|---|---|
| id | integer | DR Group ID |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
|---|---|---|
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value:<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| uuid | string | The unique identifier for the mediator service. |

partner_cluster

Partner cluster information.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

remote

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the remote cluster configuration. |
| mode | string | Specifies the mode of operation of the remote cluster. |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the remote cluster. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Initiate a switchover, heal, or switchback operation

PATCH `/cluster/metrocluster`

**Introduced In:** 9.8

Initiates a switchover, heal or switchback operation.

#### Related ONTAP commands

- `metrocluster switchover`
- `metrocluster switchback`
- `metrocluster heal`
- `metrocluster modify`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| action | string | query | False | Action to perform on the MetroCluster.<br>• enum: ["switchover", "negotiated_switchover", "negotiated_switchover_simulate", "switchback", "switchback_simulate", "heal_aggregates", "heal_root_aggregates"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| encryption_enabled | boolean | query | False | Specifies if encryption for NVLog and storage traffic is enabled.<br>• Introduced in: 9.16 |
| dr_group_id | integer | query | False | DR group ID.<br>• Introduced in: 9.16 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| configuration_type | string | Displays the MetroCluster configuration type. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| local | local | |
| remote | remote | |

**Example request**

```
{
  "configuration_type": "string",
  "local": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  },
  "remote": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  }
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2424873 | Failed to validate the node and cluster components before the "action" operation. |
| 2425138 | Switchover cannot be performed in the current DR mode. Run "metrocluster show" to view the DR mode of the local cluster, and run "switchover" only in one of the following situations: 1. The DR mode of the local cluster is "normal". 2. The DR mode of the local cluster is "partial-switchover". |
| 2425333 | Heal DR data aggregates cannot be performed in the current DR mode. Run "metrocluster show" and "metrocluster node show" to view the DR mode of the local cluster and the DR mode of the local nodes respectively, and run "heal aggregates" only in one of the following situations: 1. The DR mode of the local cluster is "switchover", and the DR mode of all local nodes is "switchover completed". 2. The DR mode of at least one local node is "heal aggrs failed". |
| 2425335 | Heal DR root aggregates cannot be performed in the current DR mode. Run "metrocluster node show" to view the DR mode of the local nodes, and run "heal root-aggregates" only in one of the following situations: 1. The DR mode of all local nodes is "heal aggrs completed". 2. The DR mode of at least one local node is "heal roots failed". |

| Error Code | Description |
| --- | --- |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2427558 | This MetroCluster operation cannot be run because another "action" operation is currently in progress. Run "metrocluster operation history show -job-id id -instance" to view the status of the currently running operation. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

dr_pairs

| Name | Type | Description |
|---|---|---|
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

cluster

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

local

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the local cluster configuration. |
| mode | string | Specifies the mode of operation of the local cluster. |
| partner_cluster_reachable | boolean | Specifies whether the partner cluster is reachable from the local cluster. |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the local cluster. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
| --- | --- | --- |
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
| --- | --- | --- |
| l3_config | l3_config | |
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

dr_group

DR group reference.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | DR Group ID |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
| --- | --- | --- |
| dr_group | dr_group | DR group reference. |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| reachable | boolean | Indicates the connectivity status of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| uuid | string | The unique identifier for the mediator service. |

partner_cluster

Partner cluster information.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

remote

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the remote cluster configuration. |
| mode | string | Specifies the mode of operation of the remote cluster. |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the remote cluster. |

metrocluster

Holds MetroCluster status and configuration parameters for the local and remote clusters. REST: /api/cluster/metrocluster

| Name | Type | Description |
|---|---|---|
| configuration_type | string | Displays the MetroCluster configuration type. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| local | local | |
| remote | remote | |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Set up a MetroCluster configuration**

POST `/cluster/metrocluster`

**Introduced In:** 9.8

Sets up a MetroCluster.

#### Required properties

- `partner_cluster.name`
- `dr_pairs`

#### Recommended optional properties

- `mediator.*`
- `mccip_ports`

#### Learn more

- DOC /cluster/metrocluster

#### Related ONTAP commands

- `metrocluster configuration-settings dr-group create`
- `metrocluster configuration-settings interface create`
- `metrocluster configuration-settings connection connect`
- `metrocluster configuration-settings mediator add`
- `storage aggregate create`
- `storage aggregate mirror`
- `metrocluster configure`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| configuration_type | string | Displays the MetroCluster configuration type. |
| dr_pairs | array[dr_pairs] | DR Pairs to create as part of a MetroCluster configure. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| local | local | |

| Name | Type | Description |
|---|---|---|
| mccip_ports | array[mccip_ports] | List of Port specifications. |
| mediator | mediator | Mediator information |
| partner_cluster | partner_cluster | Partner cluster information. |
| remote | remote | |

**Example request**

```
{
  "configuration_type": "string",
  "dr_pairs": [
    {
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "partner": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "local": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  },
  "mccip_ports": [
    {
      "l3_config": {
        "ipv4_interface": {
          "address": "10.10.10.7",
          "gateway": "10.1.1.1",
          "netmask": "24"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "string",
      "vlan_id": 200
    }
  ],
  "mediator": {
    "bluexp_account_token": "string",
    "bluexp_org_id": "string",
```

```
    "ca_certificate": "string",
    "dr_group": {
      "id": 0
    },
    "ip_address": "10.10.10.7",
    "local_mediator_connectivity": "connected",
    "password": "mypassword",
    "peer_cluster": {
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
    },
    "peer_mediator_connectivity": "connected",
    "port": 31784,
    "reachable": 1,
    "service_account_client_id": "string",
    "service_account_client_secret": "string",
    "strict_cert_validation": 1,
    "type": "string",
    "use_http_proxy_local": 1,
    "use_http_proxy_remote": 1,
    "user": "myusername",
    "uuid": "string"
  },
  "partner_cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "remote": {
    "automatic_uso_failure_domain": "string",
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
  }
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

##### Response

```
Status: 201, Created
```

##### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2432832 | Required environment variables are not set. |
| 2432833 | Operation is already running. |
| 2432834 | MetroCluster is already configured. |
| 2432835 | Operation not supported. |
| 2432836 | There are not enough disks in Pool1. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2432839 | Required parameters not set. |
| 2432840 | Configuring DR Groups |

| Error Code | Description |
| --- | --- |
| 2432841 | Generating IP addresses |
| 2432843 | Running Aggregate Recommender |
| 2432844 | Checking remote storage pool |
| 2432845 | Mirroring aggregates |
| 2432846 | Configuring MetroCluster and DR mirroring |
| 2432848 | Setting up MetroCluster |
| 2432849 | MetroCluster setup is complete |
| 2432851 | Minimum number of required data aggregates for MetroCluster configuration are still not mirrored. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

dr_pairs

| Name | Type | Description |
|---|---|---|
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

cluster

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

local

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the local cluster configuration. |
| mode | string | Specifies the mode of operation of the local cluster. |
| partner_cluster_reachable | boolean | Specifies whether the partner cluster is reachable from the local cluster. |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the local cluster. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
| --- | --- | --- |
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
| --- | --- | --- |
| l3_config | l3_config | |
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

dr_group

DR group reference.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | DR Group ID |

peer_cluster

The peer cluster that the mediator service is used for.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

mediator

Mediator information

| Name | Type | Description |
| --- | --- | --- |
| bluexp_account_token | string | BlueXP account token. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| bluexp_org_id | string | BlueXP organization ID. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| ca_certificate | string | CA certificate for ONTAP Mediator. This is optional if the certificate is already installed.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.8<br>• x-nullable: true |
| dr_group | dr_group | DR group reference. |
| ip_address | string | The IP address of the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| local_mediator_connectivity | string | Indicates the mediator connectivity status of the local cluster. Possible values are connected, unreachable, unusable and down-high-latency. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| password | string | The password used to connect to the REST server on the mediator. |
| peer_cluster | peer_cluster | The peer cluster that the mediator service is used for. |
| peer_mediator_connectivity | string | Indicates the mediator connectivity status of the peer cluster. Possible values are connected, unreachable, unknown and down-high-latency. |
| port | integer | The REST server's port number on the mediator. |
| reachable | boolean | Indicates the connectivity status of the mediator. |
| service_account_client_id | string | Client ID of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| service_account_client_secret | string | Client secret token of the BlueXP service account. This field is only applicable to the ONTAP cloud mediator.<br>• x-ntap-createOnly: true<br>• Introduced in: 9.17<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| strict_cert_validation | boolean | Indicates if strict validation of certificates is performed while making REST API calls to the mediator. This field is only applicable to the ONTAP Cloud Mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| type | string | Mediator type. This field is only applicable to the mediators in SnapMirror active sync configuration. |
| use_http_proxy_local | boolean | Indicates if the local cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• x-nullable: true |
| use_http_proxy_remote | boolean | Indicates if the remote cluster should use an http-proxy server while making REST API calls to the mediator. This field is only applicable to the ONTAP cloud mediator.<br>• Default value: 1<br>• example: true<br>• Introduced in: 9.17<br>• writeOnly: 1<br>• x-nullable: true |
| user | string | The username used to connect to the REST server on the mediator. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier for the mediator service. |

partner_cluster

Partner cluster information.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

remote

| Name | Type | Description |
| --- | --- | --- |
| automatic_uso_failure_domain | string | This parameter specifies the configuration of automatic switchover.<br>The valid values for the parameter are ':'<br>cluster':' triggers an unplanned switchover if all nodes in a DR cluster are down.<br>dr_group':' triggers an unplanned switchover if both nodes of a DR group are down.<br>disabled':' automatic switchover is disabled.<br>If the cluster is not reachable due to errors, the parameter value will be set to<br>not_reachable. This value is read only.<br>If the cluster configuration is unknown, the parameter value will be set to<br>unknown and the value is read only. |
| cluster | cluster | |
| configuration_state | string | Indicates the state of the remote cluster configuration. |
| mode | string | Specifies the mode of operation of the remote cluster. |

| Name | Type | Description |
| --- | --- | --- |
| periodic_check_enabled | boolean | Indicates whether or not a periodic check is enabled on the remote cluster. |

metrocluster

Holds MetroCluster status and configuration parameters for the local and remote clusters. REST: /api/cluster/metrocluster

| Name | Type | Description |
| --- | --- | --- |
| configuration_type | string | Displays the MetroCluster configuration type. |
| dr_pairs | array[dr_pairs] | DR Pairs to create as part of a MetroCluster configure. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| local | local | |
| mccip_ports | array[mccip_ports] | List of Port specifications. |
| mediator | mediator | Mediator information |
| partner_cluster | partner_cluster | Partner cluster information. |
| remote | remote | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Display MetroCluster diagnostics

### Display MetroCluster diagnostics

#### Overview

You can use this API to initiate a MetroCluster diagnostics operation and fetch the results of a completed diagnostics operation on a MetroCluster over IP configuration. The GET operation retrieves the results of a completed diagnostics operation for the MetroCluster over IP configuration. These can include the overall high level and details for the checks done for different components. By default, the response does not include the details. If the fields query is used in the request, the response will include the details. The POST request can be used to start a MetroCluster diagnostics operation or set up a schedule for the diagnostics to be run periodically.

#### Details

Details provide a way to view all the checks done on a component and the result of each check. The details of the checks are not included
in the response by default. In order to fetch the details, use the `fields` query parameter.

- `node.details`
- `aggregate.details`
- `cluster.details`
- `volume.details`
- `connection.details`

#### Starting a MetroCluster diagnostics operation

A new MetroCluster diagnostics operation can be started by issuing a POST to /cluster/metrocluster/diagnostics. There are no extra parameters required to initiate a diagnostics operation.

### Polling the POST job for status of diagnostics operation

After a successful POST /cluster/diagnostics operation is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The POST job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in

the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

---

#### Examples

### Running the diagnostics operation

This example shows the POST request for starting a diagnostic operation for a MetroCluster over IP configuration and the responses returned:

```
#API
/api/cluster/metrocluster/diagnostics
```

### POST Request

```
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster/diagnostics
```

### POST Response

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 17:20:53 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster/diagnostics
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "f7d3804c-fcf7-11ea-acaf-005056bb47c1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f7d3804c-fcf7-11ea-acaf-005056bb47c1"
    }
  }
}
}
```

### Monitoring the job progress

Use the link provided in the response to the POST request to fetch information for the diagnostics operation job.

**Request**

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/f7d3804c-fcf7-11ea-acaf-
005056bb47c1
```

**Job status response**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 17:21:12 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 345
Content-Type: application/hal+json
{
"uuid": "f7d3804c-fcf7-11ea-acaf-005056bb47c1",
"description": "POST /api/cluster/metrocluster/diagnostics",
"state": "running",
"message": "Checking nodes...",
"code": 2432853,
"start_time": "2020-09-22T13:20:53-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f7d3804c-fcf7-11ea-acaf-005056bb47c1"
  }
}
}
```

**Final status of the diagnostics job**

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 17:29:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 372
Content-Type: application/hal+json
{
"uuid": "f7d3804c-fcf7-11ea-acaf-005056bb47c1",
"description": "POST /api/cluster/metrocluster/diagnostics",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-09-22T13:20:53-04:00",
"end_time": "2020-09-22T13:22:04-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f7d3804c-fcf7-11ea-acaf-005056bb47c1"
  }
}
}
```

#### Retrieving the diagnostics operation

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster/diagnostics
```

#### Response

```
HTTP/1.1 202 Accepted
Date: Tue, 22 Sep 2020 18:04:28 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 1005
Content-Type: application/hal+json
{
"node": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
```

```
},
"interface": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"aggregate": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"cluster": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"connection": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"volume": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"config_replication": {
  "timestamp": "2020-09-22T13:47:01-04:00",
  "state": "ok",
  "summary": {
    "message": ""
  }
},
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/diagnostics"
```

```
  }
}
}
```

#### Retrieving check details for the node component

#### Request

```
curl -X GET https://<mgmt-
ip>/api/cluster/metrocluster/diagnostics?fields=node.details
```

#### Response

```
HTTP/1.1 200 OK
Date: Thu, 10 Feb 2022 00:05:12 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 4506
Content-Type: application/hal+json
{
"node": {
  "details": [
    {
      "node": {
        "uuid": "11111111-1111-1111-1111-111111111111",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/11111111-1111-1111-1111-
111111111111"
          }
        }
      },
      "cluster": {
        "uuid": "12121212-1212-1212-1212-121212121212",
        "name": "clusterA",
        "_links": {
          "self": {
            "href": "/api/cluster/12121212-1212-1212-1212-121212121212"
          }
        }
      },
      "timestamp": "2022-02-09T18:47:00-05:00",
```

```
      "checks": [
        {
          "name": "node_reachable",
          "result": "ok"
        },
        {
          "name": "metrocluster_ready",
          "result": "ok"
        },
        {
          "name": "local_ha_partner",
          "result": "ok"
        },
        {
          "name": "ha_mirroring_on",
          "result": "ok"
        },
        {
          "name": "ha_mirroring_op_state",
          "result": "ok"
        },
        {
          "name": "symmetric_ha_relationship",
          "result": "ok"
        },
        {
          "name": "remote_dr_partner",
          "result": "ok"
        },
        {
          "name": "dr_mirroring_on",
          "result": "ok"
        },
        {
          "name": "dr_mirroring_op_state",
          "result": "ok"
        },
        {
          "name": "symmetric_dr_relationship",
          "result": "ok"
        },
        {
          "name": "remote_dr_auxiliary_partner",
          "result": "ok"
        },
        {
```

```
          "name": "symmetric_dr_auxiliary_relationship",
          "result": "ok"
        },
        {
          "name": "storage_failover_enabled",
          "result": "ok"
        },
        {
          "name": "has_intercluster_lif",
          "result": "ok"
        },
        {
          "name": "node_object_limit",
          "result": "ok"
        },
        {
          "name": "automatic_uso",
          "result": "ok"
        }
      ]
    },
    {
      "node": {
        "uuid": "22222222-2222-2222-2222-222222222222",
        "name": "node2",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/22222222-2222-2222-2222-
222222222222"
          }
        }
      },
      "cluster": {
        "uuid": "23232323-2323-2323-2323-232323232323",
        "name": "clusterB",
        "_links": {
          "self": {
            "href": "/api/cluster/23232323-2323-2323-2323-232323232323"
          }
        }
      },
      "timestamp": "2022-02-09T18:47:00-05:00",
      "checks": [
        {
          "name": "node_reachable",
          "result": "ok"
```

```
      },
      {
        "name": "metrocluster_ready",
        "result": "ok"
      },
      {
        "name": "local_ha_partner",
        "result": "ok"
      },
      {
        "name": "ha_mirroring_on",
        "result": "ok"
      },
      {
        "name": "ha_mirroring_op_state",
        "result": "ok"
      },
      {
        "name": "symmetric_ha_relationship",
        "result": "ok"
      },
      {
        "name": "remote_dr_partner",
        "result": "ok"
      },
      {
        "name": "dr_mirroring_on",
        "result": "ok"
      },
      {
        "name": "dr_mirroring_op_state",
        "result": "ok"
      },
      {
        "name": "symmetric_dr_relationship",
        "result": "ok"
      },
      {
        "name": "remote_dr_auxiliary_partner",
        "result": "ok"
      },
      {
        "name": "symmetric_dr_auxiliary_relationship",
        "result": "ok"
      },
      {
```

```
          "name": "storage_failover_enabled",
          "result": "ok"
        },
        {
          "name": "has_intercluster_lif",
          "result": "ok"
        },
        {
          "name": "node_object_limit",
          "result": "ok"
        },
        {
          "name": "automatic_uso",
          "result": "ok"
        }
      ]
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/diagnostics"
  }
}
}
```


#### Retrieving check details for the volume component

#### Request

```
curl -X GET https://<mgmt-
ip>/api/cluster/metrocluster/diagnostics?fields=volume.details
```

#### Response

```
HTTP/1.1 200 OK
Cache-Control: no-cache,no-store,must-revalidate
Connection: close
Date: Fri, 08 Apr 2022 20:07:38 GMT
Server: libzapid-httpd
Vary: Accept-Encoding
Content-Length: 928
Content-Type: application/hal+json
{
"volume": {
  "details": [
    {
      "checks": [
        {
          "name": "unmirrored_flexgroups",
          "result": "ok",
        }
      ]
    },
    {
      "checks": [
        {
          "name": "mixed_flexgroups",
          "result": "ok",
        }
      ]
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/diagnostics"
  }
}
}
```

### Related ONTAP Commands

- `metrocluster check run`
- `metrocluster check show`
- `metrocluster check node show`
- `metrocluster check aggregate show`
- `metrocluster check cluster show`

- `metrocluster check connection show`
- `metrocluster check volume show`

### Retrieve diagnostic operation results for a MetroCluster configuration

GET `/cluster/metrocluster/diagnostics`

**Introduced In:** 9.8

Retrieves the results of a completed diagnostic operation for the MetroCluster configuration.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| aggregate | aggregate | |
| cluster | cluster | |
| config-replication | config-replication | |
| connection | connection | |
| interface | interface | |
| node | node | |
| volume | volume | |

**Example response**


```
{
  "aggregate": {
    "details": [
      {
        "aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "checks": [
          {
            "additional_info": {
              "code": "string",
              "message": "string"
            },
            "name": "mirror_status",
            "result": "string"
          }
        ],
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "timestamp": "2016-03-10 14:35:16 -0800",
        "volume": {
          "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  },
  "cluster": {
    "details": [
      {
        "aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "checks": [
          {
            "additional_info": {
              "code": "string",
              "message": "string"
            },
            "name": "mirror_status",
            "result": "string"
          }
        ],
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "timestamp": "2016-03-10 14:35:16 -0800",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  },
  "config-replication": {
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-14 14:35:16 -0800"
  },
  "connection": {
    "details": [
      {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "connections": [
          {
            "destination_address": "string",
            "partner": {
              "node": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "node1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              },
              "type": "string"
            },
            "port": "string",
            "result": "string",
            "source_address": "string",
            "state": "string"
          }
        ],
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  },
  "interface": {
    "state": "string",
    "summary": {
      "code": "string",
```

```
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  },
  "node": {
    "details": [
      {
        "aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "checks": [
          {
            "additional_info": {
              "code": "string",
              "message": "string"
            },
            "name": "mirror_status",
            "result": "string"
          }
        ],
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "timestamp": "2016-03-10 14:35:16 -0800",
        "volume": {
```

```
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  },
  "volume": {
    "details": [
      {
        "aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "checks": [
          {
            "additional_info": {
              "code": "string",
              "message": "string"
            },
            "name": "mirror_status",
            "result": "string"
          }
        ],
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
```

```
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "timestamp": "2016-03-10 14:35:16 -0800",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "state": "string",
    "summary": {
      "code": "string",
      "message": "string"
    },
    "timestamp": "2016-03-10 14:35:16 -0800"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

| Error Code | Description |
| --- | --- |
| 2426405 | MetroCluster is not configured in cluster "cluster". Run "metrocluster show" to verify both the local and the remote clusters are configured as MetroCluster, and then try the command again. |
| 2432856 | MetroCluster diagnostics result is not available. Use the REST API GET method on "/api/cluster/metrocluster/operations?type=check&fields=*" for more information. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregate

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

additional_info

Additional information or recovery steps to take on this component.

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

metrocluster_diag_check

Generic object which can be used for various components which holds details of the checks of a component.

| Name | Type | Description |
|---|---|---|
| additional_info | additional_info | Additional information or recovery steps to take on this component. |
| name | string | Name of type of diagnostic operation run for the component. |
| result | string | Result of the diagnostic operation on this component. |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

metrocluster_diag_details

| Name | Type | Description |
|---|---|---|
| aggregate | aggregate | Aggregate |
| checks | array[metrocluster_diag_check] | Collection of MetroCluster checks done for component. |
| cluster | cluster | |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | Time check was done. |
| volume | volume | |

summary

Additional information or recovery steps to take.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

aggregate

| Name | Type | Description |
| --- | --- | --- |
| details | array[metrocluster_diag_details] | Display details of the MetroCluster check for aggregates. |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |
| timestamp | string | Time of the most recent diagnostic operation for this component |

cluster

| Name | Type | Description |
| --- | --- | --- |
| details | array[metrocluster_diag_details] | Display details of the MetroCluster check for clusters. |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | Time of the most recent diagnostic operation for this component |

config-replication

| Name | Type | Description |
| --- | --- | --- |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |
| timestamp | string | Time of the most recent diagnostic operation for this component |

partner

| Name | Type | Description |
| --- | --- | --- |
| node | node | |
| type | string | |

metrocluster_diag_connection

| Name | Type | Description |
| --- | --- | --- |
| destination_address | string | |
| partner | partner | |
| port | string | |
| result | string | Result of the diagnostic operation on this component. |
| source_address | string | |
| state | string | |

metrocluster_diag_connection_details

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster | |
| connections | array[metrocluster_diag_connection] | |

| Name | Type | Description |
| --- | --- | --- |
| node | node | |

connection

| Name | Type | Description |
| --- | --- | --- |
| details | array[metrocluster_diag_connection_details] | Display details of the MetroCluster check for connections. |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |
| timestamp | string | Time of the most recent diagnostic operation for this component |

interface

| Name | Type | Description |
| --- | --- | --- |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |
| timestamp | string | Time of the most recent diagnostic operation for this component |

node

| Name | Type | Description |
| --- | --- | --- |
| details | array[metrocluster_diag_details] | Displays details of the MetroCluster check for nodes. |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |

| Name | Type | Description |
|---|---|---|
| timestamp | string | Time of the most recent diagnostic operation for this component |

volume

| Name | Type | Description |
|---|---|---|
| details | array[metrocluster_diag_details] | Display details of the MetroCluster check for volumes. |
| state | string | Status of diagnostic operation for this component. |
| summary | summary | Additional information or recovery steps to take. |
| timestamp | string | Time of the most recent diagnostic operation for this component |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Start MetroCluster diagnostics or set up a periodic diagnostic schedule**

POST `/cluster/metrocluster/diagnostics`

**Introduced In:** 9.8

Start a MetroCluster diagnostic operation or set up a schedule for the diagnostics to be run periodically.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| schedule | integer | query | False | Shows the minutes of every hour when a job runs. Setting this parameter schedules the periodic job to be run to perform MetroCluster diagnostic. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2427132 | MetroCluster is not configured on this cluster. |
| 2432833 | Operation is already running. |
| 2432852 | MetroCluster diagnostics start |
| 2432853 | MetroCluster diagnostics job scheduled |

| Error Code | Description |
| --- | --- |
| 2432854 | MetroCluster diagnostics complete |
| 2432855 | MetroCluster diagnostics operation failed. Use the REST API GET method on "/api/cluster/metrocluster/operations?type=check&fields=*" for more information. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage MetroCluster DR groups

### Manage MetroCluster DR groups

#### Overview

You can use this API to create, perform operations, and retrieve relevant information pertaining to MetroCluster DR groups. The GET operation retrieves all the DR groups in the MetroCluster over IP configuration or a DR group information specified by the DR group id. The POST request can be used to create a new DR group in the MetroCluster over IP configuration. The DELETE operation removes a DR group information specified by

the DR group id from the existing MetroCluster over IP configuration.

#### Creating a new DR group

A new DR group in MetroCluster over IP configuration is created by issuing a POST to /cluster/metrocluster/dr-groups. Parameters are provided in the body of the POST request.
This operation requires a valid MetroCluster over IP configuration. The new nodes added belong to either the local or partner cluster.

### Fields used for setting up a new DR group

The fields used for MetroCluster APIs are either required or optional and are described as follows:

### Required configuration fields

These fields are always required for any POST /cluster/dr-groups request.

- `partner_cluster.name` - Specifies the partner cluster name to which cluster peering has been established.
- `dr_pairs` - Specifies local and DR partner node pairs. Each pair uniquely identifies a DR group.

### Optional configuration fields

This field is used to set up additional MetroCluster DR configuration.

- `mccip_ports` - Specifies relevant layer 3 network configuration information for each port. These include port name, node name, IP address, gateway, and netmask. If mccip_ports is not provided, then the API automatically generates IP addresses for the ports and creates an layer 2 network configuration.

### Polling the create job

After a successful POST /cluster/metrocluster/dr-groups is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The create job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

#### Deleting a DR group using ID

A DR group in MetroCluster over IP configuration can be deleted by issuing a DELETE to /cluster/metrocluster/dr-groups/{id}. No parameters are required for the DELETE request.
The following preparation steps must be completed on the local and partner clusters before removing a DR group.

- Move all the data volumes to another DR group.
- Move all the MDV_CRS metadata volumes to another DR group.
- Delete all the MDV_aud metadata volumes that may exist in the DR group to be removed.
- Delete all the data aggregates in the DR group to be removed. Root aggregates are not deleted.
- Migrate all the data LIFs to home nodes in another DR group.
- Migrate the cluster management LIF to a home node in another DR group. Node management and inter-cluster LIFs are not migrated.

- Transfer epsilon to a node in another DR group.
  The operation is refused if the preparation steps are not completed on the local and partner clusters.

### Polling the delete job

After a successful DELETE /cluster/metrocluster/dr-groups is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The delete job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

---

#### Examples

### Creating a DR group for MetroCluster over IP configuration

This example shows the POST body when creating a DR group for MetroCluster.

```
# API
/api/cluster/metrocluster/dr-groups
```

### POST body included from file

```
dr_group_post_body.txt:
{
"partner_cluster" : {
  "name": "mcc_siteB"
},
"dr_pairs" : [
  {
    "node" : {
      "name" : "node-e"
    },
    "partner" : {
      "name" : "node-g"
    }
  }
]
}
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster/dr-groups -H
"Content-Type: application+hal/json" -d "@dr_group_post_body.txt"
```

### Inline POST body

```
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster/dr-groups -H
"Content-Type: application+hal/json" -d '{"partner_cluster" : {"name":
"mcc_siteB" }, "dr_pairs" : [{"node" : {"name" : "node-e" }, "partner" :
{"name" : "node-g" }}]}'
```

**POST Response**

```
HTTP/1.1 202 Accepted
Date: Fri, 18 Sep 2020 20:38:05 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster/dr-groups
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "5b89472e-f9e8-11ea-9c31-005056bb42f7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/5b89472e-f9e8-11ea-9c31-005056bb42f7"
    }
  }
}
}
```

**Monitoring the job progress**

Use the link provided in the response to the POST request to fetch information for the DR group job.

**Request**

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/5b89472e-f9e8-11ea-9c31-
005056bb42f7
```

**Job status response**

The following is an example of the job status response returned by the running DR group job:

```
HTTP/1.1 200 OK
Date: Fri, 18 Sep 2020 20:40:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 373
Content-Type: application/hal+json
{
"uuid": "5b89472e-f9e8-11ea-9c31-005056bb42f7",
"description": "POST /api/cluster/metrocluster/dr-groups/",
"state": "running",
"message": "Mirroring aggregates",
"code": 2432845,
"start_time": "2020-09-18T15:38:08-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/5b89472e-f9e8-11ea-9c31-005056bb42f7"
  }
}
}
```

**Final status of a successful DR Group create workflow**

When the create job completes, the 'end_time' field is populated, and the 'state' and 'message' fields report final status.

```
HTTP/1.1 200 OK
Date: Fri, 18 Sep 2020 20:43:54 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 373
Content-Type: application/hal+json
{
"uuid": "5b89472e-f9e8-11ea-9c31-005056bb42f7",
"description": "POST /api/cluster/metrocluster/dr-groups/",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-09-18T15:51:35-04:00",
"end_time": "2020-09-18T16:10:17-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/5b89472e-f9e8-11ea-9c31-005056bb42f7"
  }
}
}
```

#### Retrieving the MetroCluster DR Groups configured in the MetroCluster over IP configuration

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster/dr-groups
```

#### Response

```
HTTP/1.1 200 OK
Date: Fri, 18 Sep 2020 20:47:05 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 849
Content-Type: application/hal+json
{
"records": [
  {
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/dr-groups/1"
      }
    }
  },
  {
    "id": 2,
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/dr-groups/2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/dr-groups"
  }
}
}
```

#### Retrieving a Specific MetroCluster DR Group

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster/dr-groups/2
```

#### Response

```
HTTP/1.1 200 OK
Date: Fri, 18 Sep 2020 20:49:05 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 1049
Content-Type: application/hal+json
{
"id": 2,
"partner_cluster": {
  "name": "mcc_siteB",
  "uuid": "ea4d7114-f97f-11ea-a4bf-005056bb070a"
},
"dr_pairs": [
  {
    "node": {
      "name": "node-e",
      "uuid": "28f71e17-f988-11ea-b1dd-005056bb47e8"
    },
    "partner": {
      "name": "node-g",
      "uuid": "1af02867-f989-11ea-b86c-005056bbe97f"
    }
  },
  {
    "node": {
      "name": "node-f",
      "uuid": "b34ae3b8-f988-11ea-866b-005056bb0934"
    },
    "partner": {
      "name": "node-h",
      "uuid": "a21a2b16-f989-11ea-98d0-005056bb321d"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/dr-groups/2"
  }
}
}
```

### Deleting a MetroCluster DR Group

#### Request

```
curl -X DELETE https://<mgmt-ip>/api/cluster/metrocluster/dr-groups/{id}
```

#### Response

```
HTTP/1.1 200 OK
Date: Tue, 22 Sep 2020 03:29:01 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "c24d1083-fc83-11ea-acaf-005056bb47c1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c24d1083-fc83-11ea-acaf-005056bb47c1"
    }
  }
}
}
```

#### Monitoring the job progress

Use the link provided in the response to the DELETE request to fetch information for the delete job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/c24d1083-fc83-11ea-acaf-
005056bb47c1
```

#### Job status response

The following is an example of the job status response returned by the MetroCluster DR Group delete job.

```
HTTP/1.1 200 OK
Date: Tue, 22 Sep 2020 03:30:01 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 374
Content-Type: application/hal+json
{
"uuid": "c24d1083-fc83-11ea-acaf-005056bb47c1",
"description": "DELETE /api/cluster/metrocluster/dr-groups/2",
"state": "running",
"message": "Unconfiguring Metrocluster DR Group",
"code": 2432859,
"start_time": "2020-09-21T23:29:01-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/c24d1083-fc83-11ea-acaf-005056bb47c1"
  }
}
}
```

**Final Status of a successful MetroCluster DR Group delete workflow**

When the delete job completes, the 'end_time' field is populated, and the 'state' and 'message' fields report the final status.

```
HTTP/1.1 200 OK
Date: Tue, 22 Sep 2020 03:38:08 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 374
Content-Type: application/hal+json
{
"uuid": "c24d1083-fc83-11ea-acaf-005056bb47c1",
"description": "DELETE /api/cluster/metrocluster/dr-groups/2",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-09-21T23:29:01-04:00",
"end_time": "2020-09-21T23:36:36-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/c24d1083-fc83-11ea-acaf-005056bb47c1"
  }
}
}
```

### Retrieve all DR groups in a MetroCluster IP configuration

GET `/cluster/metrocluster/dr-groups`

**Introduced In:** 9.8

Retrieves all the DR group in the MetroCluster over IP configuration.

#### Related ONTAP commands

- `metrocluster configuration-settings dr-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | query | False | Filter by id |
| partner_cluster.name | string | query | False | Filter by partner_cluster.name |
| partner_cluster.uuid | string | query | False | Filter by partner_cluster.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| dr_pairs.partner.name | string | query | False | Filter by dr_pairs.partner.name |
| dr_pairs.partner.uuid | string | query | False | Filter by dr_pairs.partner.uuid |
| dr_pairs.node.name | string | query | False | Filter by dr_pairs.node.name |
| dr_pairs.node.uuid | string | query | False | Filter by dr_pairs.node.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of Records |
| records | array[metrocluster_dr_group] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "dr_pairs": [
        {
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "partner": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "id": 0,
      "partner_cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  }
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_pair

| Name | Type | Description |
| --- | --- | --- |
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
| --- | --- | --- |
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
| --- | --- | --- |
| l3_config | l3_config | |

| Name | Type | Description |
|---|---|---|
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

partner_cluster

Partner cluster information.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

metrocluster_dr_group

DR group information.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| dr_pairs | array[dr_pair] | |
| id | integer | DR Group ID |
| partner_cluster | partner_cluster | Partner cluster information. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a new DR group in a MetroCluster IP configuration

POST `/cluster/metrocluster/dr-groups`

**Introduced In:** 9.8

Creates a new DR group in the MetroCluster over IP configuration.

#### Required properties

- `partner_cluster.name`
- `dr_pairs`

#### Recommended optional properties

- `mccip_ports`

#### Learn more

- DOC /cluster/metrocluster/dr-groups

#### Related ONTAP commands

- `metrocluster configuration-settings dr-group create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| dr_pairs | array[dr_pair] | |
| id | integer | DR Group ID |
| mccip_ports | array[mccip_ports] | List of Port specifications. |

| Name | Type | Description |
| --- | --- | --- |
| partner_cluster | partner_cluster | Partner cluster information. |

**Example request**

```
{
  "dr_pairs": [
    {
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "partner": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "id": 0,
  "mccip_ports": [
    {
      "l3_config": {
        "ipv4_interface": {
          "address": "10.10.10.7",
          "gateway": "10.1.1.1",
          "netmask": "24"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "string",
      "vlan_id": 200
    }
  ],
  "partner_cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2430034 | Node "node" is not found in partner cluster "cluster". Use the "cluster show" command on the partner cluster to see the list of valid node names. |
| 2430060 | The DR group cannot be created while node "node" on cluster "cluster" is configured for MetroCluster. |
| 2432833 | Operation is already running. |

| Error Code | Description |
| --- | --- |
| 2432836 | There are not enough disks in Pool1. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2432840 | Configuring DR Groups |
| 2432841 | Generating IP addresses |
| 2432844 | Checking remote storage pool |
| 2432845 | Mirroring aggregates |
| 2432846 | Configuring MetroCluster and DR mirroring |
| 2432857 | Adding new MetroCluster DR Group |
| 2432858 | MetroCluster DR Group setup done |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

dr_pair

| Name | Type | Description |
|---|---|---|
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
| --- | --- | --- |
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
| --- | --- | --- |
| l3_config | l3_config | |
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

partner_cluster

Partner cluster information.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

metrocluster_dr_group

DR group information.

| Name | Type | Description |
|---|---|---|
| dr_pairs | array[dr_pair] | |
| id | integer | DR Group ID |
| mccip_ports | array[mccip_ports] | List of Port specifications. |
| partner_cluster | partner_cluster | Partner cluster information. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Remove a DR group from a MetroCluster IP configuration**

DELETE `/cluster/metrocluster/dr-groups/{id}`

**Introduced In:** 9.8

Remove the DR group from the current MetroCluster over IP configuration specified by the DR group id.

#### Related ONTAP commands

- `metrocluster configuration-settings dr-group delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | string | path | True | |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2425574 | Two or more DR groups must be configured to remove a DR group from the MetroCluster configuration. |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2432833 | Operation is already running. |
| 2432859 | Unconfiguring MetroCluster DR Group |
| 2432860 | Unmirroring Aggregates |
| 2432861 | Unassigning Remote Disks |
| 2432862 | Disabling Cluster HA and Storage Failover HA |
| 2432863 | Disconnecting and deleting network connections |

| Error Code | Description |
| --- | --- |
| 2432864 | Unconfiguring and deleting the DR Group |
| 2432865 | Deleting MetroCluster DR Group |
| 2432866 | MetroCluster DR Group delete done |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve DR group information using the DR group ID

GET `/cluster/metrocluster/dr-groups/{id}`

**Introduced In:** 9.8

Retrieves the DR group information specified by the DR group id.

#### Related ONTAP commands

- `metrocluster configuration-settings dr-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | string | path | True | |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| dr_pairs | array[dr_pair] | |
| id | integer | DR Group ID |
| partner_cluster | partner_cluster | Partner cluster information. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "dr_pairs": [
    {
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "partner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "id": 0,
  "partner_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Local node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

partner

Partner node of the DR Group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_pair

| Name | Type | Description |
|---|---|---|
| node | node | Local node of the DR Group. |
| partner | partner | Partner node of the DR Group. |

ipv4_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

l3_config

| Name | Type | Description |
| --- | --- | --- |
| ipv4_interface | ipv4_interface | Object to setup an interface along with its default router. |

node

Node information

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

mccip_ports

Port configuration specification.
l3_config information is only needed when configuring a MetroCluster IP for use in a layer 3 network.

| Name | Type | Description |
| --- | --- | --- |
| l3_config | l3_config | |
| name | string | Port name |
| node | node | Node information |
| uuid | string | Port UUID |
| vlan_id | integer | VLAN ID |

partner_cluster

Partner cluster information.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update MetroCluster interconnects

### View and update MetroCluster interconnects

#### Overview

You can use this API to retrieve and display relevant information pertaining to MetroCluster interconnect status. The `/cluster/metrocluster/interconnects` endpoint returns a list of all the interconnects in MetroCluster and their status. Each individual interconnect can be queried individually using the `/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{adapter}` endpoint. You can also use this API to modify relevant information related to MetroCluster interconnect. These include address, netmask, and gateway. Modify a MetroCluster interconnect using the `/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{adapter}` endpoint.

---

### Examples

#### Retrieving MetroCluster interconnects

```
GET https://<mgmt-ip>/api/cluster/metrocluster/interconnects
{
  "records": [
      {
          "node": {
              "name": "cluster1_01",
              "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
              "_links": {
                  "self": {
                      "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                  }
              }
          },
          "partner_type": "ha",
          "adapter": "e0f",
          "_links": {
              "self": {
                  "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/ha/e0f"
              }
          }
      },
      {
          "node": {
              "name": "cluster1_01",
              "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
              "_links": {
                  "self": {
                      "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                  }
              }
          },
          "partner_type": "ha",
          "adapter": "e0g",
          "_links": {
              "self": {
                  "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/ha/e0g"
              }
          }
```

```
      },
      {
          "node": {
              "name": "cluster1_01",
              "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
              "_links": {
                  "self": {
                      "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                  }
              }
          },
          "partner_type": "dr",
          "adapter": "e0f",
          "_links": {
              "self": {
                  "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/dr/e0f"
              }
          }
      },
      {
          "node": {
              "name": "cluster1_01",
              "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
              "_links": {
                  "self": {
                      "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                  }
              }
          },
          "partner_type": "dr",
          "adapter": "e0g",
          "_links": {
              "self": {
                  "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/dr/e0g"
              }
          }
      },
      {
          "node": {
              "name": "cluster1_01",
```

```
                "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                    }
                }
            },
            "partner_type": "aux",
            "adapter": "e0f",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/aux/e0f"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_01",
                "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                    }
                }
            },
            "partner_type": "aux",
            "adapter": "e0g",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/aux/e0g"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
```

```
00505682dc8b"
                    }
                }
            },
            "partner_type": "ha",
            "adapter": "e0f",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/ha/e0f"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                    }
                }
            },
            "partner_type": "ha",
            "adapter": "e0g",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/ha/e0g"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                    }
                }
            },
```

```
            "partner_type": "dr",
            "adapter": "e0f",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/dr/e0f"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                    }
                }
            },
            "partner_type": "dr",
            "adapter": "e0g",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/dr/e0g"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                    }
                }
            },
            "partner_type": "aux",
            "adapter": "e0f",
            "_links": {
                "self": {
```

```
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/aux/e0f"
                }
            }
        },
        {
            "node": {
                "name": "cluster1_02",
                "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                    }
                }
            },
            "partner_type": "aux",
            "adapter": "e0g",
            "_links": {
                "self": {
                    "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/aux/e0g"
                }
            }
        }
    ],
    "num_records": 12,
    "_links": {
        "self": {
            "href": "/api/cluster/metrocluster/interconnects"
        }
    }
}
```


**Retrieves information about a specific MetroCluster interconnect**

```
https://<mgmt-ip>/api/cluster/metrocluster/interconnects/774b4fbc-86f9-
11e9-9051-005056825c71/aux/e0f
{
  "node": {
      "name": "cluster1_01",
      "uuid": "46147363-9857-11e9-9a55-005056828eb9",
      "_links": {
          "self": {
              "href": "/api/cluster/nodes/46147363-9857-11e9-9a55-
005056828eb9"
          }
      }
  },
  "partner_type": "aux",
  "adapter": "e0f",
  "state": "up",
  "type": "iwarp",
  "interfaces": [
    {
      "address": "10.2.3.5",
      "netmask": "255.255.255.0"
    }
  ],
  "mirror": {
      "state": "online",
      "enabled": true
  },
  "multipath_policy": "static_map",
  "_links": {
      "self": {
          "href": "/api/cluster/metrocluster/interconnects/46147363-9857-
11e9-9a55-005056828eb9/ha/e0f"
      }
  }
}
```

**This example shows how to modify the network address assigned to the home port. Fields required: address.**

```
curl -X PATCH https://<mgmt-
ip>/api/cluster/metrocluster/interconnects/3e1bfd38-ffd2-11eb-bcb7-
005056aceaa9/ha/e0g -d '{"interfaces": [{"address": "1.2.3.4"}]}'
```

**PATCH Response**

```
HTTP/1.1 200 OK
Cache-Control: no-cache,no-store,must-revalidate
Connection: close
Date: Fri, 20 Aug 2021 21:58:36 GMT
Server: libzapid-httpd
Content-Length: 3
Content-Type: application/hal+json
X-Content-Type-Options: nosniff
{
}
```

**This example shows how to modify the netmask assigned to the interface. Be sure to change to a valid subnet. Fields required: netmask.**

```
curl -X PATCH https://<mgmt-
ip>/api/cluster/metrocluster/interconnects/3e1bfd38-ffd2-11eb-bcb7-
005056aceaa9/ha/e0g -d '{"interfaces": [{"netmask": "2.2.2.2"}]}'
```

**PATCH Response**

```
HTTP/1.1 200 OK
Cache-Control: no-cache,no-store,must-revalidate
Connection: close
Date: Fri, 20 Aug 2021 22:11:35 GMT
Server: libzapid-httpd
Content-Length: 3
Content-Type: application/hal+json
X-Content-Type-Options: nosniff
{
}
```

**This example shows how to modify the gateway assigned to the interface. Please make sure to update it on the switch/router first. Assuming it is a new one, the new gateway and IP address must reside in the same subnet range as the interface IP address. Fields required: gateway.**

```
curl -X PATCH https://<mgmt-
ip>/api/cluster/metrocluster/interconnects/3e1bfd38-ffd2-11eb-bcb7-
005056aceaa9/ha/e0g -d '{"interfaces": [{"gateway": "1.2.3.4"}]}'
```

### PATCH Response

```
HTTP/1.1 200 OK
Cache-Control: no-cache,no-store,must-revalidate
Connection: close
Date: Fri, 20 Aug 2021 22:11:35 GMT
Server: libzapid-httpd
Content-Length: 3
Content-Type: application/hal+json
X-Content-Type-Options: nosniff
{
}
```

### Retrieve interconnect adapter information for nodes in MetroCluster

GET `/cluster/metrocluster/interconnects`

**Introduced In:** 9.8

Retrieves a list of interconnect adapter information for nodes in the MetroCluster.

#### Related ONTAP Commands

- `metrocluster interconnect show`

#### Learn more

- DOC /cluster/metrocluster/interconnects

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| adapter | string | query | False | Filter by adapter |
| state | string | query | False | Filter by state |
| multipath_policy | string | query | False | Filter by multipath_policy<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| vlan_id | integer | query | False | Filter by vlan_id<br>• Introduced in: 9.9<br>• Max value: 4095<br>• Min value: 1 |
| interfaces.netmask | string | query | False | Filter by interfaces.netmask<br>• Introduced in: 9.9 |
| interfaces.address | string | query | False | Filter by interfaces.address<br>• Introduced in: 9.9 |
| interfaces.gateway | string | query | False | Filter by interfaces.gateway<br>• Introduced in: 9.9 |
| type | string | query | False | Filter by type |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| partner_type | string | query | False | Filter by partner_type |
| mirror.enabled | boolean | query | False | Filter by mirror.enabled<br>• Introduced in: 9.11 |
| mirror.state | string | query | False | Filter by mirror.state<br>• Introduced in: 9.11 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[metrocluster_interconnect] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "adapter": "string",
      "interfaces": [
        {
          "address": "10.10.10.7",
          "gateway": "10.1.1.1",
          "netmask": "24"
        }
      ],
      "mirror": {
        "state": "string"
      },
      "multipath_policy": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "partner_type": "string",
      "state": "string",
      "type": "string",
      "vlan_id": 0
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2424994 | MetroCluster is not supported when HAOSC is not set to "mcc", "mcc-2n" or "mccip". Run the "ha-config modify chassis" and "ha-config modify controller" commands in maintenance mode to change the HAOSC settings on node "node" in cluster "cluster". |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2427132 | MetroCluster is not configured on this cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

interfaces

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

mirror

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies the administrative state of the NVRAM mirror between partner nodes. |
| state | string | Specifies the operational state of the NVRAM mirror between partner nodes. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

metrocluster_interconnect

Data for a MetroCluster interconnect. REST: /api/cluster/metrocluster/interconnects

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| adapter | string | Adapter |
| interfaces | array[interfaces] | List of objects which contain interface information such as its IP address, netmask and gateway. |
| mirror | mirror | |
| multipath_policy | string | Displays the NVRAM mirror multipath policy for the nodes configured in a MetroCluster. |
| node | node | |
| partner_type | string | Partner type |
| state | string | Adapter status |
| type | string | Adapter type |
| vlan_id | integer | VLAN ID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve information about a MetroCluster interconnect for a partner type and adapter

GET `/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{adapter}`

**Introduced In:** 9.8

Retrieves information about a MetroCluster Interconnect for a specific partner type and adapter.

#### Related ONTAP Commands

- `metrocluster interconnect show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| partner_type | string | path | True | DR Partner type |
| adapter | string | path | True | Interconnect adapter. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | Adapter |
| interfaces | array[interfaces] | List of objects which contain interface information such as its IP address, netmask and gateway. |
| mirror | mirror | |
| multipath_policy | string | Displays the NVRAM mirror multipath policy for the nodes configured in a MetroCluster. |
| node | node | |
| partner_type | string | Partner type |
| state | string | Adapter status |
| type | string | Adapter type |
| vlan_id | integer | VLAN ID |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adapter": "string",
  "interfaces": [
    {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  ],
  "mirror": {
    "state": "string"
  },
  "multipath_policy": "string",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "partner_type": "string",
  "state": "string",
  "type": "string",
  "vlan_id": 0
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2427132 | MetroCluster is not configured on this cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

interfaces

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

mirror

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies the administrative state of the NVRAM mirror between partner nodes. |
| state | string | Specifies the operational state of the NVRAM mirror between partner nodes. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a MetroCluster interconnect interface

PATCH `/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{adapter}`

**Introduced In:** 9.10

Updates a MetroCluster interconnect interface.

#### Related ONTAP commands

- `metrocluster configuration-settings interface modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| partner_type | string | path | True | DR Partner type |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| adapter | string | path | True | Interconnect adapter |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| adapter | string | Adapter |
| interfaces | array[interfaces] | List of objects which contain interface information such as its IP address, netmask and gateway. |
| mirror | mirror | |

| Name | Type | Description |
|---|---|---|
| multipath_policy | string | Displays the NVRAM mirror multipath policy for the nodes configured in a MetroCluster. |
| node | node | |
| partner_type | string | Partner type |
| state | string | Adapter status |
| type | string | Adapter type |
| vlan_id | integer | VLAN ID |

**Example request**

```
{
  "adapter": "string",
  "interfaces": [
    {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  ],
  "mirror": {
    "state": "string"
  },
  "multipath_policy": "string",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "partner_type": "string",
  "state": "string",
  "type": "string",
  "vlan_id": 0
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

**Example response**

```
{
  "job": {
    "uuid": "string"
  }
}
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2425734 | An internal error has occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |
| 2430347 | The interface is already configured for port "adapter" with IP address "interfaces.address" on node "node.name". To configure a new IP address, run the "metrocluster configuration-settings interface delete" command to delete this interface and then run the "metrocluster configuration-settings interface create" command with the new IP address. |
| 2430348 | The interface is already configured for port "adapter" with gateway "interfaces.gateway" on node "node.name". To configure a new gateway or no gateway, run the "metrocluster configuration-settings interface delete" command to delete this interface and then run the "metrocluster configuration-settings interface create" command with the new gateway or no gateway. |

| Error Code | Description |
| --- | --- |
| 2430352 | The interface is already configured for port "adapter" with netmask "interfaces.netmask" on node "node.name". To configure a new netmask, run the "metrocluster configuration-settings interface delete" command to delete this interface and then run the "metrocluster configuration-settings interface create" command with the new netmask. |
| 2430445 | Configuration for the interface "port" on node "node.name" in cluster "cluster" does not exist. Use the "metrocluster configuration-settings interface show" command to list the valid interfaces configured for MetroCluster over IP. |
| 2432885 | Unable to modify the interface for adapter "port". Reason: At least one interface parameter is required to modify the interface. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

interfaces

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

mirror

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies the administrative state of the NVRAM mirror between partner nodes. |
| state | string | Specifies the operational state of the NVRAM mirror between partner nodes. |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

metrocluster_interconnect

Data for a MetroCluster interconnect. REST: /api/cluster/metrocluster/interconnects

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | Adapter |
| interfaces | array[interfaces] | List of objects which contain interface information such as its IP address, netmask and gateway. |
| mirror | mirror | |
| multipath_policy | string | Displays the NVRAM mirror multipath policy for the nodes configured in a MetroCluster. |
| node | node | |
| partner_type | string | Partner type |
| state | string | Adapter status |
| type | string | Adapter type |
| vlan_id | integer | VLAN ID |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve MetroCluster node configurations

### Retrieve MetroCluster node configurations

#### Overview

Retrieves the configuration information for the nodes in the MetroCluster configuration.

---

### Example

```
GET https://<mgmt-ip>/api/cluster/metrocluster/nodes
{
  "records": [
      {
          "dr_group_id": 1,
          "cluster": {
              "name": "cluster1",
              "uuid": "8f77de32-9857-11e9-9a55-005056828eb9",
              "_links": {
                  "self": {
                      "href": "/api/cluster"
                  }
              }
          },
          "node": {
              "name": "cluster1_01",
              "uuid": "46147363-9857-11e9-9a55-005056828eb9",
              "_links": {
                  "self": {
                      "href": "/api/cluster/nodes/46147363-9857-11e9-9a55-
005056828eb9"
                  }
              }
          },
          "dr_mirroring_state": "enabled",
          "configuration_state": "configured",
          "_links": {
```

```
                "self": {
                    "href": "/api/cluster/metrocluster/nodes/46147363-9857-
11e9-9a55-005056828eb9"
                }
            }
        },
        {
            "dr_group_id": 1,
            "cluster": {
                "name": "cluster1",
                "uuid": "8f77de32-9857-11e9-9a55-005056828eb9",
                "_links": {
                    "self": {
                        "href": "/api/cluster"
                    }
                }
            },
            "node": {
                "name": "cluster1_02",
                "uuid": "cf1dc67f-9857-11e9-bf80-005056829db6",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/cf1dc67f-9857-11e9-bf80-
005056829db6"
                    }
                }
            },
            "dr_mirroring_state": "enabled",
            "configuration_state": "configured",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/nodes/cf1dc67f-9857-
11e9-bf80-005056829db6"
                }
            }
        },
        {
            "dr_group_id": 1,
            "cluster": {
                "name": "cluster3",
                "uuid": "aa8aa15a-9857-11e9-80c9-00505682e684",
                "_links": {
                    "self": {
                        "href": "/api/cluster/peers/aa8aa15a-9857-11e9-80c9-
00505682e684/cluster"
                    }
```

```
                }
            },
            "node": {
                "name": "cluster3_01",
                "uuid": "5b3b983b-9857-11e9-80c9-00505682e684",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/5b3b983b-9857-11e9-80c9-
00505682e684"
                    }
                }
            },
            "dr_mirroring_state": "enabled",
            "configuration_state": "configured",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/nodes/5b3b983b-9857-
11e9-80c9-00505682e684"
                }
            }
        },
        {
            "dr_group_id": 1,
            "cluster": {
                "name": "cluster3",
                "uuid": "aa8aa15a-9857-11e9-80c9-00505682e684",
                "_links": {
                    "self": {
                        "href": "/api/cluster/peers/aa8aa15a-9857-11e9-80c9-
00505682e684/cluster"
                    }
                }
            },
            "node": {
                "name": "cluster3_02",
                "uuid": "45bff538-9858-11e9-a624-005056820377",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/45bff538-9858-11e9-a624-
005056820377"
                    }
                }
            },
            "dr_mirroring_state": "enabled",
            "configuration_state": "configured",
            "_links": {
```

```
                "self": {
                    "href": "/api/cluster/metrocluster/nodes/45bff538-9858-
11e9-a624-005056820377"
                }
            }
        }
    ],
    "num_records": 4,
    "_links": {
        "self": {
            "href": "/api/cluster/metrocluster/nodes?fields=%2A"
        }
    }
}
```


---

### Retrieve MetroCluster nodes and configurations

GET `/cluster/metrocluster/nodes`

**Introduced In:** 9.8

Retrieves MetroCluster nodes and their configurations.

#### Related ONTAP Commands

- `metrocluster node show`

#### Learn more

- DOC /cluster/metrocluster/nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| is_mccip | boolean | query | False | Filter by is_mccip<br>• Introduced in: 9.12 |
| dr_group_id | integer | query | False | Filter by dr_group_id |
| dr_partner.uuid | string | query | False | Filter by dr_partner.uuid<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| dr_partner.system_id | string | query | False | Filter by dr_partner.system_id<br>• Introduced in: 9.12 |
| dr_partner.name | string | query | False | Filter by dr_partner.name<br>• Introduced in: 9.12 |
| dr_mirroring_state | string | query | False | Filter by dr_mirroring_state |
| ha_partner.uuid | string | query | False | Filter by ha_partner.uuid<br>• Introduced in: 9.12 |
| ha_partner.system_id | string | query | False | Filter by ha_partner.system_id<br>• Introduced in: 9.12 |
| ha_partner.name | string | query | False | Filter by ha_partner.name<br>• Introduced in: 9.12 |
| encryption_enabled | boolean | query | False | Filter by encryption_enabled<br>• Introduced in: 9.15 |
| dr_partner_cluster.name | string | query | False | Filter by dr_partner_cluster.name<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| dr_partner_cluster.uuid | string | query | False | Filter by dr_partner_cluster.uuid<br>• Introduced in: 9.12 |
| limit_enforcement | string | query | False | Filter by limit_enforcement<br>• Introduced in: 9.12 |
| configuration_state | string | query | False | Filter by configuration_state |
| node.uuid | string | query | False | Filter by node.uuid |
| node.system_id | string | query | False | Filter by node.system_id<br>• Introduced in: 9.12 |
| node.name | string | query | False | Filter by node.name |
| dr_auxiliary_partner.uuid | string | query | False | Filter by dr_auxiliary_partner.uuid<br>• Introduced in: 9.12 |
| dr_auxiliary_partner.system_id | string | query | False | Filter by dr_auxiliary_partner.system_id<br>• Introduced in: 9.12 |
| dr_auxiliary_partner.name | string | query | False | Filter by dr_auxiliary_partner.name<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| automatic_uso | boolean | query | False | Filter by automatic_uso<br><br>• Introduced in: 9.9 |
| dr_auxiliary_cluster.name | string | query | False | Filter by dr_auxiliary_cluster.name<br><br>• Introduced in: 9.12 |
| dr_auxiliary_cluster.uuid | string | query | False | Filter by dr_auxiliary_cluster.uuid<br><br>• Introduced in: 9.12 |
| cluster.name | string | query | False | Filter by cluster.name |
| cluster.uuid | string | query | False | Filter by cluster.uuid |
| ha_partner_cluster.name | string | query | False | Filter by ha_partner_cluster.name<br><br>• Introduced in: 9.12 |
| ha_partner_cluster.uuid | string | query | False | Filter by ha_partner_cluster.uuid<br><br>• Introduced in: 9.12 |
| dr_operation_state | string | query | False | Filter by dr_operation_state<br><br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[metrocluster_node] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "configuration_state": "string",
      "dr_auxiliary_cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "dr_auxiliary_partner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "system_id": "string",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "dr_group_id": 0,
      "dr_mirroring_state": "string",
      "dr_operation_state": "string",
      "dr_partner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "system_id": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "dr_partner_cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "ha_partner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "system_id": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "ha_partner_cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "limit_enforcement": "string",
      "node": {
        "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "system_id": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_auxiliary_cluster

DR AUX cluster.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_auxiliary_partner

AUX partner node.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

dr_partner

DR partner node.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

dr_partner_cluster

DR partner cluster.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

ha_partner

HA partner node.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

ha_partner_cluster

HA partner cluster.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

metrocluster_node

Data for a node in a MetroCluster. REST: /api/cluster/metrocluster/nodes

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| automatic_uso | boolean | Specifies if automatic unplanned switchover is enabled. |
| cluster | cluster | |
| configuration_state | string | Configuration state of the node. |
| dr_auxiliary_cluster | dr_auxiliary_cluster | DR AUX cluster. |
| dr_auxiliary_partner | dr_auxiliary_partner | AUX partner node. |
| dr_group_id | integer | DR Group ID. |
| dr_mirroring_state | string | State of the DR mirroring configuration. |
| dr_operation_state | string | State of the DR operation. |
| dr_partner | dr_partner | DR partner node. |
| dr_partner_cluster | dr_partner_cluster | DR partner cluster. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| ha_partner | ha_partner | HA partner node. |

| Name | Type | Description |
|---|---|---|
| ha_partner_cluster | ha_partner_cluster | HA partner cluster. |
| is_mccip | boolean | Indicates whether the configuration type is MCC-IP. |
| limit_enforcement | string | Indicates if the node object limits are enforced. |
| node | node | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the node configuration in MetroCluster

GET `/cluster/metrocluster/nodes/{node.uuid}`

**Introduced In:** 9.8

Retrieves the node configuration in the MetroCluster.

#### Related ONTAP Commands

- `metrocluster node show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| automatic_uso | boolean | Specifies if automatic unplanned switchover is enabled. |
| cluster | cluster | |
| configuration_state | string | Configuration state of the node. |
| dr_auxiliary_cluster | dr_auxiliary_cluster | DR AUX cluster. |
| dr_auxiliary_partner | dr_auxiliary_partner | AUX partner node. |
| dr_group_id | integer | DR Group ID. |
| dr_mirroring_state | string | State of the DR mirroring configuration. |
| dr_operation_state | string | State of the DR operation. |
| dr_partner | dr_partner | DR partner node. |
| dr_partner_cluster | dr_partner_cluster | DR partner cluster. |
| encryption_enabled | boolean | Indicates if the encryption for NVLog and storage traffic is enabled. |
| ha_partner | ha_partner | HA partner node. |
| ha_partner_cluster | ha_partner_cluster | HA partner cluster. |

| Name | Type | Description |
| --- | --- | --- |
| is_mccip | boolean | Indicates whether the configuration type is MCC-IP. |
| limit_enforcement | string | Indicates if the node object limits are enforced. |
| node | node | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "configuration_state": "string",
  "dr_auxiliary_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "dr_auxiliary_partner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "system_id": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "dr_group_id": 0,
  "dr_mirroring_state": "string",
  "dr_operation_state": "string",
  "dr_partner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
```

```
    "name": "node1",
    "system_id": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "dr_partner_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "ha_partner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "system_id": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "ha_partner_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "limit_enforcement": "string",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "system_id": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_auxiliary_cluster

DR AUX cluster.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

dr_auxiliary_partner

AUX partner node.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

dr_partner

DR partner node.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

dr_partner_cluster

DR partner cluster.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

ha_partner

HA partner node.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | |
| system_id | string | |
| uuid | string | |

ha_partner_cluster

HA partner cluster.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

node

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | |
| system_id | string | |

| Name | Type | Description |
|---|---|---|
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve MetroCluster operations

### Retrieve MetroCluster operations

### Overview

Retrieves a list of recent MetroCluster operations. To view more information about a specific operation, use the `/cluster/metrocluster/operations/{uuid}` API endpoint.

---

#### Examples

### Retrieves all MetroCluster operations

```
GET https://<mgmt-ip>/api/cluster/metrocluster/operations?fields=*
{
  "records": [
      {
          "uuid": "a14ae39f-8d85-11e9-b4a7-00505682dc8b",
          "type": "check",
          "state": "successful",
```

```
            "start_time": "2019-06-14T11:15:00-07:00",
            "end_time": "2019-06-14T11:16:08-07:00",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/operations/a14ae39f-
8d85-11e9-b4a7-00505682dc8b"
                }
            }
        },
        {
            "uuid": "7058df27-8d85-11e9-bbc9-005056826931",
            "type": "configure",
            "state": "successful",
            "start_time": "2019-06-12T19:46:27-07:00",
            "end_time": "2019-06-12T19:48:17-07:00",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/operations/7058df27-
8d85-11e9-bbc9-005056826931"
                }
            }
        },
        {
            "uuid": "7849515d-8d84-11e9-bbc9-005056826931",
            "type": "connect",
            "state": "successful",
            "start_time": "2019-06-12T19:39:30-07:00",
            "end_time": "2019-06-12T19:42:02-07:00",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/operations/7849515d-
8d84-11e9-bbc9-005056826931"
                }
            }
        },
        {
            "uuid": "331c79ad-8d84-11e9-b4a7-00505682dc8b",
            "type": "interface_create",
            "state": "successful",
            "start_time": "2019-06-12T19:37:35-07:00",
            "end_time": "2019-06-12T19:37:41-07:00",
            "_links": {
                "self": {
                    "href": "/api/cluster/metrocluster/operations/331c79ad-
8d84-11e9-b4a7-00505682dc8b"
                }
```

```
            }
        }
    ],
    "num_records": 4,
    "_links": {
        "self": {
            "href": "/api/cluster/metrocluster/operations?fields=%2A"
        }
    }
}
```

#### Retrieves Information about a specific MetroCluster operation

```
GET https://<mgmt-ip>/api/cluster/metrocluster/operations/0db12274-86fd-
11e9-8053-00505682c342
{
  "uuid": "0db12274-86fd-11e9-8053-00505682c342",
  "name": "check",
  "state": "successful",
  "start_time": "2019-06-06T16:15:01-07:00",
  "end_time": "2019-06-06T16:16:05-07:00",
  "_links": {
      "self": {
          "href": "/api/cluster/metrocluster/operations/0db12274-86fd-
11e9-8053-00505682c342"
      }
  }
}
```

---

### Retrieve MetroCluster operations on the local cluster

GET `/cluster/metrocluster/operations`

**Introduced In:** 9.8

Retrieves the list of MetroCluster operations on the local cluster.

#### Related ONTAP Commands

- `metrocluster operation history show`

#### Learn more

- DOC /cluster/metrocluster/operations

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| end_time | string | query | False | Filter by end_time |
| uuid | string | query | False | Filter by uuid |
| errors | string | query | False | Filter by errors |
| command_line | string | query | False | Filter by command_line |
| state | string | query | False | Filter by state |
| additional_info | string | query | False | Filter by additional_info |
| start_time | string | query | False | Filter by start_time |
| type | string | query | False | Filter by type |
| node.name | string | query | False | Filter by node.name<br>• Introduced in: 9.9 |
| node.uuid | string | query | False | Filter by node.uuid<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[metrocluster_operation] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "additional_info": "MetroCluster switchover with auto heal
completed successfully.",
      "command_line": "metrocluster switchover",
      "end_time": "2016-03-10 14:35:16 -0800",
      "errors": [
        "siteB (warning): Unable to prepare the partner cluster for a
pending switchback operation. Reason: entry doesn't exist. Reboot the
nodes in the partner cluster before using the \"metrocluster
switchback\" command."
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "start_time": "2016-03-10 14:33:16 -0800",
      "state": "completed_with_warnings",
      "type": "switchover",
      "uuid": "11111111-2222-3333-4444-abcdefabcdef"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Node from where the command is executed.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

metrocluster_operation

Data for a MetroCluster operation. REST: /api/cluster/metrocluster/operations

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| additional_info | string | Additional information for the auto heal. |
| command_line | string | Command line executed with the options specified. |

| Name | Type | Description |
|---|---|---|
| end_time | string | End Time |
| errors | array[string] | List of errors in the operation. |
| node | node | Node from where the command is executed. |
| start_time | string | Start Time |
| state | string | Indicates the state of the operation. |
| type | string | Name of the operation. |
| uuid | string | Identifier for the operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve information about a MetroCluster operation**

GET `/cluster/metrocluster/operations/{uuid}`

**Introduced In:** 9.8

Retrieves information about a specific MetroCluster operation.

#### Related ONTAP Commands

- `metrocluster operation show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Unique identifier for the operation. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| additional_info | string | Additional information for the auto heal. |
| command_line | string | Command line executed with the options specified. |
| end_time | string | End Time |
| errors | array[string] | List of errors in the operation. |
| node | node | Node from where the command is executed. |
| start_time | string | Start Time |
| state | string | Indicates the state of the operation. |
| type | string | Name of the operation. |
| uuid | string | Identifier for the operation. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "additional_info": "MetroCluster switchover with auto heal completed
successfully.",
  "command_line": "metrocluster switchover",
  "end_time": "2016-03-10 14:35:16 -0800",
  "errors": [
    "siteB (warning): Unable to prepare the partner cluster for a
pending switchback operation. Reason: entry doesn't exist. Reboot the
nodes in the partner cluster before using the \"metrocluster
switchback\" command."
  ],
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "start_time": "2016-03-10 14:33:16 -0800",
  "state": "completed_with_warnings",
  "type": "switchover",
  "uuid": "11111111-2222-3333-4444-abcdefabcdef"
}
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2425734 | An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Node from where the command is executed.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Retrieve MetroCluster SVMs

## Retrieve MetroCluster SVMs

### Overview

Retrieves configuration information for all pairs of SVMs in MetroCluster.

---

### Related ONTAP commands

- `metrocluster vserver show`

#### Examples

```
GET https://<mgmt-ip>/api/cluster/metrocluster/svms/?fields=*"
{
"records": [
  {
    "cluster": {
      "uuid": "9623385a-6c4c-11ec-a8cc-005056aca0c8",
      "name": "cluster1"
    },
    "svm": {
      "uuid": "2ea76ca7-6c5f-11ec-b430-005056aca0c8",
      "name": "svm1"
    },
    "partner_svm": {
      "uuid": "2ea76ca7-6c5f-11ec-b430-005056aca0c8",
      "name": "svm1-mc"
    },
    "configuration_state": "healthy",
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/svms/9623385a-6c4c-11ec-a8cc-
005056aca0c8/2ea76ca7-6c5f-11ec-b430-005056aca0c8"
      }
    }
  },
  {
    "cluster": {
      "uuid": "988d33a0-6c4c-11ec-8e28-005056aceeed",
      "name": "cluster2"
    },
    "svm": {
      "uuid": "2fa16461-6c5f-11ec-8f69-005056aceeed",
      "name": "svm2"
```

```
    },
    "partner_svm": {
      "uuid": "2fa16461-6c5f-11ec-8f69-005056aceeed",
      "name": "svm2-mc"
    },
    "configuration_state": "healthy",
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/svms/988d33a0-6c4c-11ec-8e28-
005056aceeed/2fa16461-6c5f-11ec-8f69-005056aceeed"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/svms/?fields=*"
  }
}
}
```


---

**Retrieve configuration information for all SVM pairs in a MetroCluster configuration**

GET `/cluster/metrocluster/svms`

**Introduced In:** 9.11

Retrieves configuration information for all pairs of SVMs in MetroCluster. REST /api/cluster/metrocluster/svms/?

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | query | False | Filter by svm.uuid |
| svm.name | string | query | False | Filter by svm.name |
| cluster.name | string | query | False | Filter by cluster.name |
| cluster.uuid | string | query | False | Filter by cluster.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| configuration_state | string | query | False | Filter by configuration_state |
| failed_reason | string | query | False | Filter by failed_reason |
| partner_svm.name | string | query | False | Filter by partner_svm.name |
| partner_svm.uuid | string | query | False | Filter by partner_svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[metrocluster_svm] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "configuration_state": "string",
      "failed_reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "partner_svm": {
        "name": "string",
        "uuid": "string"
      },
      "svm": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

Reason for SVM object replication failure.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |

partner_svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | MetroCluster partner SVM name. |
| uuid | string | MetroCluster partner SVM UUID. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

metrocluster_svm

Retrieves configuration information for all pairs of SVMs in MetroCluster. REST /api/cluster/metrocluster/svms/?

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| cluster | cluster | |
| configuration_state | string | Configuration state. |
| failed_reason | error | Reason for SVM object replication failure. |
| partner_svm | partner_svm | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve configuration information for an SVM

GET `/cluster/metrocluster/svms/{cluster.uuid}/{svm.uuid}`

**Introduced In:** 9.11

Retrieves configuration information for an SVM in a MetroCluster relationship.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster.uuid | string | path | True | Cluster ID |
| svm.uuid | string | path | True | SVM UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| cluster | cluster | |
| configuration_state | string | Configuration state. |
| failed_reason | error | Reason for SVM object replication failure. |
| partner_svm | partner_svm | |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "configuration_state": "string",
  "failed_reason": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "partner_svm": {
    "name": "string",
    "uuid": "string"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

Reason for SVM object replication failure.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

partner_svm

| Name | Type | Description |
|---|---|---|
| name | string | MetroCluster partner SVM name. |
| uuid | string | MetroCluster partner SVM UUID. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster nodes

### Manage cluster nodes

### Overview

You can use this API to add nodes to a cluster, update node-specific configurations, and retrieve the current node configuration details.

#### Adding a node to a cluster

You can add a node to a cluster by issuing a POST /cluster/nodes request to a node currently in the cluster. All nodes must be running the same version of ONTAP to use this API. Mixed version joins are not supported in this release. You can provide properties as fields in the body of the POST request to configure node-specific

settings. On a successful request, POST /cluster/nodes returns a status code of 202 and job information in the body of the request. You can use the /cluster/jobs APIs to track the status of the node add job.

### Fields used for adding a node

Fields used for the /cluster/nodes APIs fall into the following categories:

- Required node fields
- Optional fields
- Network interface fields
- Records field

### Required node fields

The following field is required for any POST /cluster/nodes request:

- cluster_interface.ip.address

### Optional fields

All of the following fields are used to set up additional cluster-wide configurations:

- name
- location
- records

### Network interface fields

You can set a node-specific configuration for each node by using the POST /cluster/nodes API. If you provide a field in the body of a node, provide it for all nodes in the POST body.
You can provide the node management interface for each node if all node management interfaces in the cluster use the same subnet mask. If the node management interfaces use different subnet masks, use the /network/ip/interfaces API to configure the node management interfaces.

### The records field

To add multiple nodes to the cluster in one request, provide an array named "records" with multiple node entries. Each node entry in "records" must follow the required and optional fields listed previously. When only adding a single node, you do not need a "records" field. See "Examples" for an example of how to use the "records" field.

### Create recommended aggregates parameter

When you set the "create_recommended_aggregates" parameter to "true", aggregates based on an optimal layout recommended by the system are created on each of the nodes being added to the cluster. The default setting is "false".

---

#### Modifying node configurations

The following fields can be used to modify a node configuration:

- name
- location

---

#### Modifying service processor configurations

When modifying the "service_processor" properties, the job returns success immediately if valid network information is passed in. The values remain in their old state until the network information changes have taken effect on the service processor. You can poll the modified properties until the values are updated.

---

#### Deleting a node from a cluster

You can delete a node from the cluster. Before deleting a node from the cluster, shut down all of the node's shared resources, such as virtual interfaces to clients. If any of the node's shared resources are still active, the command fails.
You can use the "force" flag to forcibly remove a node that is down and cannot be brought online to remove its shared resources. This flag is set to "false" by default.

---

#### Node state

The node "state" field in the /cluster/nodes API represents the current operational state of individual nodes. Note that the state of a node is a transient value and can change depending on the current condition of the node, especially during reboot, takeover, and giveback.
Possible values for the node state are:

- *up* - Node is fully operational and is able to accept and handle management requests. It is connected to a majority of healthy (up) nodes in the cluster through the cluster interconnect and all critical services are online.
- *booting* - Node is starting up and is not yet fully functional. It might not yet be accessible through the management interface or cluster interconnect. One or more critical services are offline on the node and the node is not taken over. The HA partner reports the node's firmware state as "SF_BOOTING", "SF_BOOTED", or "SF_CLUSTERWAIT".
- *down* - Node is known to be down. It cannot be reached through the management interface or cluster interconnect. The HA partner can be reached and reports that the node is halted/rebooted without takeover. Or, the HA partner cannot be reached (or no SFO configured) but the node shutdown request has been recorded by the quorum change coordinator. The state is reported by the node's HA partner.
- *taken_over* - Node is taken over by its HA partner. The state is reported by the node's HA partner.
- *waiting_for_giveback* - Node is taken over by its HA partner and is now ready and waiting for giveback. To bring the node up, either issue the "giveback" command to the HA partner node or wait for auto-giveback, if enabled. The state is reported by the node's HA partner.
- *degraded* - Node is known to be up but is not yet fully functional. The node can be reached through the cluster interconnect but one or more critical services are offline. Or, the node is not reachable but the node's HA partner can be reached and reports that the node is up with firmware state "SF_UP".
- *unknown* - Node state cannot be determined.

---

#### HA

The "ha" field in the /cluster/nodes API shows the takeover and giveback states of the node along with the current values of the HA fields "enabled" and "auto_giveback".
You can modify the HA fields "enabled" and "auto_giveback", which will change the HA states of the node.

### Takeover

The takeover "state" field shows the different takeover states of the node. When the state is "failed", the "code" and "message" fields display.
Possible values for takeover states are:

- *not_attempted* - Takeover operation is not started and takeover is possible.
- *not_possible* - Takeover operation is not possible. Check the failure message.
- *in_progress* - Takeover operation is in progress. The node is taking over its partner.
- *in_takeover* - Takeover operation is complete.
- *failed* - Takeover operation failed. Check the failure message.

Possible values for takeover failure code and messages are:

- *code*: 852130
  *message*: Failed to initiate takeover. Run the "storage failover show-takeover" command for more information.
- *code*: 852131
  *message*: Takeover cannot be completed. Reason: disabled.

### Giveback

The giveback "state" field shows the different giveback states of the node. When the state is "failed", the "code" and "message" fields display.
Possible values for giveback states are:

- *nothing_to_giveback* - Node does not have partner aggregates to giveback.
- *not_attempted* - Giveback operation is not started.
- *in_progress* - Giveback operation is in progress.
- *failed* - Giveback operation failed. Check the failure message.

Possible values for giveback failure codes and messages are:

- *code*: 852126
  *message*: Failed to initiate giveback. Run the "storage failover show-giveback" command for more information.

---

#### Performance monitoring

Performance of a node can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the performance of a node in terms of cpu utilization. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

The following examples show how to add nodes to a cluster, update node properties, shutdown and reboot a node, and remove a node from the cluster.

### Adding a single node with a minimal configuration

```
# Body
add_single_node.txt(body):
{
"cluster_interface": {
  "ip": {
    "address": "1.1.1.1"
  }
}
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/nodes" -d
"@add_single_node.txt"
```

### Adding multiple nodes in the same request and creating recommended aggregates

```
# Body
add_multiple_nodes.txt(body):
{
"records": [
    {
        "name": "node1",
        "cluster_interface": {
          "ip": {
            "address": "1.1.1.1"
          }
        }
    },
    {
        "name": "node2",
        "cluster_interface": {
          "ip": {
            "address": "2.2.2.2"
          }
        }
    }
]
}

# Request
curl -X POST "https://<mgmt-
ip>/api/cluster/nodes?create_recommended_aggregates=true" -d
"@add_multiple_nodes.txt"
```

#### Modifying a cluster-wide configuration

```
# Body
modify_name_and_location.txt(body):
{
"name": "renamedNode",
"location": "newLocation"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes" -d
"@modify_name_and_location.txt"
```

#### Shutting down a node

```
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes/{uuid}?action=shutdown"
```

#### Powering off a node using SP assistance

```
curl -X PATCH "https://<mgmt-
ip>/api/cluster/nodes/{uuid}?action=power_off"
```

---

#### Deleting a node from a cluster

```
curl -X DELETE "https://<mgmt-ip>/api/cluster/nodes/{uuid}"
```

#### Force a node deletion from a cluster

```
curl -X DELETE "https://<mgmt-ip>/api/cluster/nodes/{uuid}?force=true"
```

---

#### Retrieving the state of all nodes in a cluster

```
#Request
curl -siku admin -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=state"

#Response
{
"records": [
  {
    "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
    "name": "node2",
    "state": "up",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/54440ec3-6127-11e9-a959-005056bb76f9"
      }
    }
  },
  {
    "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
    "name": "node1",
    "state": "up",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e02dbef1-6126-11e9-b8fb-005056bb9ce4"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=state"
  }
}
}
```

### Retrieving nodes that are in the spare low condition in a cluster

```
# Request
curl -siku admin -X GET "https://<mgmt-
ip>/api/cluster/nodes?fields=is_spares_low"

#Response
{
"records": [
  {
    "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
    "name": "node2",
    "spares_low": true,
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/54440ec3-6127-11e9-a959-005056bb76f9"
      }
    }
  },
  {
    "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
    "name": "node1",
    "spares_low": false,
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e02dbef1-6126-11e9-b8fb-005056bb9ce4"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=state"
  }
}
}
```

---

#### Retrieving statistics and metric for a node

In this example, the API returns the "statistics" and "metric" properties.

```
#Request
curl -siku admin -X GET "https://<mgmt-
ip>/api/cluster/nodes?fields=statistics,metric"

#Response
{
    "records": [
    {
      "uuid": "6b29327b-21ca-11ea-99aa-005056bb420b",
      "name": "example_node_name",
      "metric": {
        "timestamp": "2019-12-19T15:50:45Z",
        "duration": "PT15S",
        "status": "ok",
        "processor_utilization": 3
      },
      "statistics": {
        "timestamp": "2019-12-19T15:50:48Z",
        "status": "ok",
        "processor_utilization_raw": 6409411622,
        "processor_utilization_base": 74330229886
      }
    }
  ],
  "num_records": 1
}
```

#### Retrieving takeover and giveback failure codes and messages

```
#Request
curl -siku admin -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=ha"

#Response
{
"records": [
  {
    "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
    "name": "node2",
    "ha": {
      "enabled": false,
      "auto_giveback": false,
      "partners": [
        {
```

```
          "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
          "name": "node1"
        }
      ],
      "giveback": {
            "state": "nothing_to_giveback"
      },
      "takeover": {
        "state": "not_possible",
        "failure": {
          "message": "Takeover cannot be completed. Reason: disabled.",
          "code": 852131
        }
      },
      "ports": [
        {
          "name": "e0h"
        },
        {
          "name": "N/A"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/54440ec3-6127-11e9-a959-005056bb76f9"
      }
    }
  },
  {
    "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
    "name": "node1",
    "ha": {
      "enabled": false,
      "auto_giveback": false,
      "partners": [
        {
          "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
          "name": "node2"
        }
      ],
      "giveback": {
            "state": "nothing_to_giveback"
      },
      "takeover": {
        "state": "not_possible",
```

```
        "failure": {
          "message": "Takeover cannot be completed. Reason: disabled.",
          "code": 852131
        }
      },
      "ports": [
        {
          "name": "e0h"
        },
        {
          "name": "N/A"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e02dbef1-6126-11e9-b8fb-005056bb9ce4"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=state"
  }
}
}
```


---

#### Retrieving external cache information for a node

In this example, the API returns the external_cache property.

```
#Request
curl -siku admin -X GET "https://<mgmt-
ip>/api/cluster/nodes?fields=external_cache"

#Response
{
"records": [
  {
    "uuid": "71af8235-bea9-11eb-874a-005056bbab13",
    "name": "node2",
    "external_cache": {
      "is_enabled": false,
      "is_hya_enabled": true,
      "is_rewarm_enabled": false,
      "pcs_size": 256
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/71af8235-bea9-11eb-874a-005056bbab13"
      }
    }
  },
  {
    "uuid": "8c4cbf08-bea9-11eb-b8ae-005056bb16aa",
    "name": "node1",
    "external_cache": {
      "is_enabled": false,
      "is_hya_enabled": true,
      "is_rewarm_enabled": false,
      "pcs_size": 256
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/8c4cbf08-bea9-11eb-b8ae-005056bb16aa"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=external_cache"
  }
}
}
```

### Retrieve the nodes in a cluster

GET `/cluster/nodes`

**Introduced In:** 9.6

Retrieves the nodes in the cluster.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `system node show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ha.interconnect.adapter | string | query | False | Filter by ha.interconnect.adapter<br>• Introduced in: 9.11 |
| ha.interconnect.state | string | query | False | Filter by ha.interconnect.state<br>• Introduced in: 9.11 |
| ha.takeover_check.takeover_possible | boolean | query | False | Filter by ha.takeover_check.takeover_possible<br>• Introduced in: 9.14 |
| ha.takeover_check.reasons | string | query | False | Filter by ha.takeover_check.reasons<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ha.ports.number | integer | query | False | Filter by ha.ports.number<br><br>• Introduced in: 9.7 |
| ha.ports.state | string | query | False | Filter by ha.ports.state<br><br>• Introduced in: 9.7 |
| ha.auto_giveback | boolean | query | False | Filter by ha.auto_giveback |
| ha.partners.name | string | query | False | Filter by ha.partners.name |
| ha.partners.uuid | string | query | False | Filter by ha.partners.uuid |
| ha.takeover.failure.code | integer | query | False | Filter by ha.takeover.failure.code<br><br>• Introduced in: 9.7 |
| ha.takeover.failure.message | string | query | False | Filter by ha.takeover.failure.message<br><br>• Introduced in: 9.7 |
| ha.takeover.state | string | query | False | Filter by ha.takeover.state<br><br>• Introduced in: 9.7 |
| ha.type | string | query | False | Filter by ha.type<br><br>• Introduced in: 9.17 |
| ha.enabled | boolean | query | False | Filter by ha.enabled |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ha.giveback.failure.message | string | query | False | Filter by ha.giveback.failure.message<br><br>• Introduced in: 9.7 |
| ha.giveback.failure.code | integer | query | False | Filter by ha.giveback.failure.code<br><br>• Introduced in: 9.7 |
| ha.giveback.status.aggregate.name | string | query | False | Filter by ha.giveback.status.aggregate.name<br><br>• Introduced in: 9.11 |
| ha.giveback.status.aggregate.uuid | string | query | False | Filter by ha.giveback.status.aggregate.uuid<br><br>• Introduced in: 9.11 |
| ha.giveback.status.error.code | string | query | False | Filter by ha.giveback.status.error.code<br><br>• Introduced in: 9.11 |
| ha.giveback.status.error.message | string | query | False | Filter by ha.giveback.status.error.message<br><br>• Introduced in: 9.11 |
| ha.giveback.status.state | string | query | False | Filter by ha.giveback.status.state<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ha.giveback.state | string | query | False | Filter by ha.giveback.state<br>• Introduced in: 9.7 |
| membership | string | query | False | Filter by membership |
| version.major | integer | query | False | Filter by version.major |
| version.full | string | query | False | Filter by version.full |
| version.generation | integer | query | False | Filter by version.generation |
| version.minor | integer | query | False | Filter by version.minor |
| external_cache.is_rewarm_enabled | boolean | query | False | Filter by external_cache.is_rewarm_enabled<br>• Introduced in: 9.10 |
| external_cache.is_enabled | boolean | query | False | Filter by external_cache.is_enabled<br>• Introduced in: 9.10 |
| external_cache.pcs_size | integer | query | False | Filter by external_cache.pcs_size<br>• Introduced in: 9.10 |
| external_cache.is_hya_enabled | boolean | query | False | Filter by external_cache.is_hya_enabled<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| system_id | string | query | False | Filter by system_id<br>• Introduced in: 9.7 |
| serial_number | string | query | False | Filter by serial_number |
| management_interfaces.ip.address | string | query | False | Filter by management_interfaces.ip.address |
| management_interfaces.name | string | query | False | Filter by management_interfaces.name |
| management_interfaces.uuid | string | query | False | Filter by management_interfaces.uuid |
| owner | string | query | False | Filter by owner<br>• Introduced in: 9.9 |
| uuid | string | query | False | Filter by uuid |
| anti_ransomware_version | string | query | False | Filter by anti_ransomware_version<br>• Introduced in: 9.16 |
| hw_assist.status.partner.state | string | query | False | Filter by hw_assist.status.partner.state<br>• Introduced in: 9.11 |
| hw_assist.status.partner.port | integer | query | False | Filter by hw_assist.status.partner.port<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| hw_assist.status.partner.ip | string | query | False | Filter by hw_assist.status.partner.ip<br><br>• Introduced in: 9.11 |
| hw_assist.status.local.state | string | query | False | Filter by hw_assist.status.local.state<br><br>• Introduced in: 9.11 |
| hw_assist.status.local.port | integer | query | False | Filter by hw_assist.status.local.port<br><br>• Introduced in: 9.11 |
| hw_assist.status.local.ip | string | query | False | Filter by hw_assist.status.local.ip<br><br>• Introduced in: 9.11 |
| hw_assist.status.enabled | boolean | query | False | Filter by hw_assist.status.enabled<br><br>• Introduced in: 9.11 |
| nvlog.backing_type | string | query | False | Filter by nvlog.backing_type<br><br>• Introduced in: 9.17 |
| nvlog.swap_mode | string | query | False | Filter by nvlog.swap_mode<br><br>• Introduced in: 9.17 |
| model | string | query | False | Filter by model |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| vendor_serial_number | string | query | False | Filter by vendor_serial_number<br><br>• Introduced in: 9.7 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.8 |
| statistics.processor_utilization_raw | integer | query | False | Filter by statistics.processor_utilization_raw<br><br>• Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.8 |
| statistics.processor_utilization_base | integer | query | False | Filter by statistics.processor_utilization_base<br><br>• Introduced in: 9.8 |
| controller.board | string | query | False | Filter by controller.board<br><br>• Introduced in: 9.9 |
| controller.failed_power_supply.count | integer | query | False | Filter by controller.failed_power_supply.count<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| controller.failed_power_supply.message.code | string | query | False | Filter by controller.failed_power_supply.message.code<br><br>• Introduced in: 9.9 |
| controller.failed_power_supply.message.message | string | query | False | Filter by controller.failed_power_supply.message.message<br><br>• Introduced in: 9.9 |
| controller.cpu.count | integer | query | False | Filter by controller.cpu.count<br><br>• Introduced in: 9.9 |
| controller.cpu.firmware_release | string | query | False | Filter by controller.cpu.firmware_release<br><br>• Introduced in: 9.9 |
| controller.cpu.processor | string | query | False | Filter by controller.cpu.processor<br><br>• Introduced in: 9.9 |
| controller.frus.type | string | query | False | Filter by controller.frus.type |
| controller.frus.id | string | query | False | Filter by controller.frus.id |
| controller.frus.state | string | query | False | Filter by controller.frus.state |
| controller.flash_cache.slot | string | query | False | Filter by controller.flash_cache.slot |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| controller.flash_cache.part_number | string | query | False | Filter by controller.flash_cache.part_number |
| controller.flash_cache.serial_number | string | query | False | Filter by controller.flash_cache.serial_number |
| controller.flash_cache.firmware_file | string | query | False | Filter by controller.flash_cache.firmware_file<br>• Introduced in: 9.9 |
| controller.flash_cache.model | string | query | False | Filter by controller.flash_cache.model |
| controller.flash_cache.state | string | query | False | Filter by controller.flash_cache.state |
| controller.flash_cache.device_id | integer | query | False | Filter by controller.flash_cache.device_id<br>• Introduced in: 9.9 |
| controller.flash_cache.capacity | integer | query | False | Filter by controller.flash_cache.capacity |
| controller.flash_cache.firmware_version | string | query | False | Filter by controller.flash_cache.firmware_version |
| controller.flash_cache.hardware_revision | string | query | False | Filter by controller.flash_cache.hardware_revision |
| controller.over_temperature | string | query | False | Filter by controller.over_temperature |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| controller.failed_fan.count | integer | query | False | Filter by controller.failed_fan.count<br><br>• Introduced in: 9.9 |
| controller.failed_fan.message.code | string | query | False | Filter by controller.failed_fan.message.code<br><br>• Introduced in: 9.9 |
| controller.failed_fan.message.message | string | query | False | Filter by controller.failed_fan.message.message<br><br>• Introduced in: 9.9 |
| controller.memory_size | integer | query | False | Filter by controller.memory_size<br><br>• Introduced in: 9.9 |
| location | string | query | False | Filter by location |
| cluster_interfaces.ip.address | string | query | False | Filter by cluster_interfaces.ip.address |
| cluster_interfaces.name | string | query | False | Filter by cluster_interfaces.name |
| cluster_interfaces.uuid | string | query | False | Filter by cluster_interfaces.uuid |
| nvram.id | integer | query | False | Filter by nvram.id<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| nvram.battery_state | string | query | False | Filter by nvram.battery_state<br><br>• Introduced in: 9.9 |
| external_cache_bypass.enabled | boolean | query | False | Filter by external_cache_bypass.enabled<br><br>• Introduced in: 9.17 |
| external_cache_bypass.large_read_ops_allow_percent | integer | query | False | Filter by external_cache_bypass.large_read_ops_allow_percent<br><br>• Introduced in: 9.17 |
| system_aggregate.name | string | query | False | Filter by system_aggregate.name<br><br>• Introduced in: 9.14 |
| system_aggregate.uuid | string | query | False | Filter by system_aggregate.uuid<br><br>• Introduced in: 9.14 |
| service_processor.auto_config.ipv4_subnet | string | query | False | Filter by service_processor.auto_config.ipv4_subnet<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| service_processor.auto_config.ipv6_subnet | string | query | False | Filter by service_processor.auto_config.ipv6_subnet<br><br>• Introduced in: 9.11 |
| service_processor.web_service.enabled | boolean | query | False | Filter by service_processor.web_service.enabled<br><br>• Introduced in: 9.14 |
| service_processor.web_service.limit_access | boolean | query | False | Filter by service_processor.web_service.limit_access<br><br>• Introduced in: 9.14 |
| service_processor.type | string | query | False | Filter by service_processor.type<br><br>• Introduced in: 9.10 |
| service_processor.mac_address | string | query | False | Filter by service_processor.mac_address |
| service_processor.ssh_info.allowed_addresses | string | query | False | Filter by service_processor.ssh_info.allowed_addresses<br><br>• Introduced in: 9.10 |
| service_processor.state | string | query | False | Filter by service_processor.state |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| service_processor.ipv6_interface.is_ipv6_ra_enabled | boolean | query | False | Filter by service_processor.ipv6_interface.is_ipv6_ra_enabled<br>• Introduced in: 9.14 |
| service_processor.ipv6_interface.setup_state | string | query | False | Filter by service_processor.ipv6_interface.setup_state<br>• Introduced in: 9.14 |
| service_processor.ipv6_interface.gateway | string | query | False | Filter by service_processor.ipv6_interface.gateway |
| service_processor.ipv6_interface.enabled | boolean | query | False | Filter by service_processor.ipv6_interface.enabled<br>• Introduced in: 9.14 |
| service_processor.ipv6_interface.router_ip | string | query | False | Filter by service_processor.ipv6_interface.router_ip<br>• Introduced in: 9.14 |
| service_processor.ipv6_interface.link_local_ip | string | query | False | Filter by service_processor.ipv6_interface.link_local_ip<br>• Introduced in: 9.14 |
| service_processor.ipv6_interface.netmask | integer | query | False | Filter by service_processor.ipv6_interface.netmask |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| service_processor.ipv6_interface.address | string | query | False | Filter by service_processor.ipv6_interface.address |
| service_processor.autoupdate_enabled | boolean | query | False | Filter by service_processor.autoupdate_enabled<br>• Introduced in: 9.10 |
| service_processor.is_ip_configured | boolean | query | False | Filter by service_processor.is_ip_configured<br>• Introduced in: 9.10 |
| service_processor.ipv4_interface.gateway | string | query | False | Filter by service_processor.ipv4_interface.gateway |
| service_processor.ipv4_interface.address | string | query | False | Filter by service_processor.ipv4_interface.address |
| service_processor.ipv4_interface.netmask | string | query | False | Filter by service_processor.ipv4_interface.netmask |
| service_processor.ipv4_interface.setup_state | string | query | False | Filter by service_processor.ipv4_interface.setup_state<br>• Introduced in: 9.14 |
| service_processor.ipv4_interface.enabled | boolean | query | False | Filter by service_processor.ipv4_interface.enabled<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| service_processor.dhcp_enabled | boolean | query | False | Filter by service_processor.dhcp_enabled |
| service_processor.api_service.enabled | boolean | query | False | Filter by service_processor.api_service.enabled<br><br>• Introduced in: 9.11 |
| service_processor.api_service.limit_access | boolean | query | False | Filter by service_processor.api_service.limit_access<br><br>• Introduced in: 9.11 |
| service_processor.api_service.port | integer | query | False | Filter by service_processor.api_service.port<br><br>• Introduced in: 9.11 |
| service_processor.primary.is_current | boolean | query | False | Filter by service_processor.primary.is_current<br><br>• Introduced in: 9.10 |
| service_processor.primary.version | string | query | False | Filter by service_processor.primary.version<br><br>• Introduced in: 9.10 |
| service_processor.primary.state | string | query | False | Filter by service_processor.primary.state<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| service_processor.last_update_state | string | query | False | Filter by service_processor.last_update_state<br>• Introduced in: 9.10 |
| service_processor.firmware_version | string | query | False | Filter by service_processor.firmware_version |
| service_processor.backup.version | string | query | False | Filter by service_processor.backup.version<br>• Introduced in: 9.10 |
| service_processor.backup.is_current | boolean | query | False | Filter by service_processor.backup.is_current<br>• Introduced in: 9.10 |
| service_processor.backup.state | string | query | False | Filter by service_processor.backup.state<br>• Introduced in: 9.10 |
| service_processor.link_status | string | query | False | Filter by service_processor.link_status |
| vm.provider_type | string | query | False | Filter by vm.provider_type<br>• Introduced in: 9.7 |
| storage_configuration | string | query | False | Filter by storage_configuration<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| system_machine_type | string | query | False | Filter by system_machine_type<br>• Introduced in: 9.7 |
| uptime | integer | query | False | Filter by uptime |
| metrocluster.type | string | query | False | Filter by metrocluster.type<br>• Introduced in: 9.8 |
| metrocluster.custom_vlan_capable | boolean | query | False | Filter by metrocluster.custom_vlan_capable<br>• Introduced in: 9.8 |
| metrocluster.ports.name | string | query | False | Filter by metrocluster.ports.name<br>• Introduced in: 9.8 |
| date | string | query | False | Filter by date |
| snaplock.compliance_clock_time | string | query | False | Filter by snaplock.compliance_clock_time<br>• Introduced in: 9.12 |
| metric.uuid | string | query | False | Filter by metric.uuid<br>• Introduced in: 9.10 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.8 |
| metric.processor_utilization | integer | query | False | Filter by metric.processor_utilization<br><br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.8 |
| state | string | query | False | Filter by state<br><br>• Introduced in: 9.7 |
| name | string | query | False | Filter by name |
| is_spares_low | boolean | query | False | Filter by is_spares_low<br><br>• Introduced in: 9.10 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anti_ransomware_version": "1.0",
      "cluster_interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "controller": {
        "board": "System Board XXVIII",
        "cpu": {
          "count": 20,
          "firmware_release": "string",
          "processor": "string"
        },
        "failed_fan": {
          "count": 1,
          "message": {
            "code": "111411207",
            "message": "There are no failed fans."
```

```
          }
        },
        "failed_power_supply": {
          "count": 1,
          "message": {
            "code": "111411208",
            "message": "There are no failed power supplies."
          }
        },
        "flash_cache": [
          {
            "capacity": 1024000000000,
            "device_id": 0,
            "firmware_file": "X9170_O000Z6300NVM",
            "firmware_version": "NA05",
            "hardware_revision": "A1",
            "model": "X1970A",
            "part_number": "119-00207",
            "serial_number": "A22P5061550000187",
            "slot": "6-1",
            "state": "string"
          }
        ],
        "frus": [
          {
            "id": "string",
            "state": "string",
            "type": "string"
          }
        ],
        "memory_size": 1024000000,
        "over_temperature": "string"
      },
      "date": "2019-04-17 11:49:26 -0400",
      "external_cache": {
        "is_enabled": 1,
        "is_hya_enabled": 1,
        "is_rewarm_enabled": 1
      },
      "external_cache_bypass": {
        "enabled": 1,
        "large_read_ops_allow_percent": 100,
        "reset": 1
      },
      "ha": {
        "giveback": {
```

```
          "failure": {
            "code": 852126,
            "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
          },
          "state": "failed",
          "status": [
            {
              "aggregate": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "aggr1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              },
              "error": {
                "code": "852126",
                "message": "string"
              },
              "state": "string"
            }
          ]
        },
        "interconnect": {
          "adapter": "MVIA-RDMA",
          "state": "string"
        },
        "partners": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "ports": [
          {
            "number": 0,
            "state": "active"
          }
        ],
```

```
        "takeover": {
          "failure": {
            "code": 852130,
            "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
          },
          "state": "failed"
        },
        "takeover_check": {
          "reasons": [
            "string"
          ]
        },
        "type": "string"
      },
      "hw_assist": {
        "status": {
          "local": {
            "ip": "string",
            "state": "string"
          },
          "partner": {
            "ip": "string",
            "state": "string"
          }
        }
      },
      "location": "rack 2 row 5",
      "management_interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "membership": "string",
      "metric": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "processor_utilization": 13,
        "status": "ok",
        "timestamp": "2017-01-25 11:20:13 +0000",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "metrocluster": {
        "ports": [
          {
            "name": "e1b"
          }
        ],
        "type": "string"
      },
      "model": "FAS3070",
      "name": "node-01",
      "nvlog": {
        "backing_type": "string",
        "swap_mode": "string"
      },
      "nvram": {
        "battery_state": "string",
        "id": 0
      },
      "owner": "Example Corp",
      "serial_number": "4048820-60-9",
      "service_processor": {
        "api_service": {
          "port": 0
        },
        "auto_config": {
          "ipv4_subnet": "ipv4_mgmt",
          "ipv6_subnet": "ipv6_mgmt"
        },
        "backup": {
          "state": "string",
          "version": "11.6"
        },
        "firmware_version": "string",
        "ipv4_interface": {
          "address": "10.10.10.7",
          "gateway": "10.1.1.1",
          "netmask": "255.255.0.0",
```

```
          "setup_state": "string"
        },
        "ipv6_interface": {
          "address": "fd20:8b1e:b255:5011:10:141:4:97",
          "gateway": "fd20:8b1e:b255:5011:10::1",
          "link_local_ip": "FE80::/10",
          "netmask": 64,
          "router_ip": "2001:0db8:85a3:0000:0000:8a2e:0370:7334",
          "setup_state": "string"
        },
        "last_update_state": "string",
        "link_status": "string",
        "mac_address": "string",
        "primary": {
          "state": "string",
          "version": "11.6"
        },
        "ssh_info": {
          "allowed_addresses": [
            "10.10.10.7/24"
          ]
        },
        "state": "string",
        "type": "string"
      },
      "snaplock": {
        "compliance_clock_time": "2018-06-04 19:00:00 +0000"
      },
      "state": "string",
      "statistics": {
        "processor_utilization_base": 12345123,
        "processor_utilization_raw": 13,
        "status": "ok",
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "storage_configuration": "string",
      "system_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "system_id": 92027651,
```

```
      "system_machine_type": "7Y56-CTOWW1",
      "uptime": 300536,
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
      "vendor_serial_number": 791603000068,
      "version": {
        "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
        "generation": 9,
        "major": 4,
        "minor": 0
      },
      "vm": {
        "provider_type": "string"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
| --- | --- | --- |
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
| --- | --- | --- |
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
|---|---|---|
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
|---|---|---|
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
| --- | --- | --- |
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
| --- | --- | --- |
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
| --- | --- | --- |
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |
| local | local | |
| partner | partner | |

hw_assist

The hardware assist information.

| Name | Type | Description |
|---|---|---|
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
|---|---|---|
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |

| Name | Type | Description |
| --- | --- | --- |
| autoupdate_enabled | boolean | Indicates whether the service processor can be automatically updated from ONTAP.<br>• Introduced in: 9.10<br>• x-ntap-readModify: true<br>• x-nullable: true |
| backup | backup | Provides the properties of the service processor backup partition. |
| dhcp_enabled | boolean | Set to "true" to use DHCP to configure an IPv4 interface. Do not provide values for address, netmask and gateway when set to "true". |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| ipv6_interface | ipv6_interface | Object to setup an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |

| Name | Type | Description |
| --- | --- | --- |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |

| Name | Type | Description |
| --- | --- | --- |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Cloud provider where the VM is hosted. |

records

Complete node information

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |

| Name | Type | Description |
| --- | --- | --- |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| hw_assist | hw_assist | The hardware assist information. |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metric | metric | CPU performance for the nodes. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| statistics | statistics | Raw CPU performance for the nodes. |

| Name | Type | Description |
| --- | --- | --- |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |
| vm | vm | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add a node or nodes to a cluster

POST `/cluster/nodes`

**Introduced In:** 9.6

Adds a node or nodes to the cluster.

#### Required properties

- `cluster_interface.ip.address`

#### Related ONTAP commands

- `cluster add-node`
- `network interface create`
- `storage aggregate auto-provision`
- `system node modify`
- `system service-processor network modify`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| create_recommended_aggregates | boolean | query | False | Creates aggregates based on an optimal layout recommended by the system.<br>• Default value:<br>• Introduced in: 9.7 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interface | cluster_interface | The cluster network IP address of the node to be added. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interface | management_interface | The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes. |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |

| Name | Type | Description |
|---|---|---|
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

**Example request**


```
{
  "anti_ransomware_version": "1.0",
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "cluster_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "board": "System Board XXVIII",
    "cpu": {
      "count": 20,
      "firmware_release": "string",
      "processor": "string"
    },
    "failed_fan": {
      "count": 1,
      "message": {
        "code": "111411207",
        "message": "There are no failed fans."
      }
    },
    "failed_power_supply": {
      "count": 1,
      "message": {
        "code": "111411208",
        "message": "There are no failed power supplies."
      }
    },
    "flash_cache": [
      {
        "capacity": 1024000000000,
        "device_id": 0,
        "firmware_file": "X9170_O000Z6300NVM",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
```

```
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": "string",
        "state": "string",
        "type": "string"
      }
    ],
    "memory_size": 1024000000,
    "over_temperature": "string"
  },
  "date": "2019-04-17 11:49:26 -0400",
  "external_cache": {
    "is_enabled": 1,
    "is_hya_enabled": 1,
    "is_rewarm_enabled": 1
  },
  "external_cache_bypass": {
    "enabled": 1,
    "large_read_ops_allow_percent": 100,
    "reset": 1
  },
  "ha": {
    "giveback": {
      "failure": {
        "code": 852126,
        "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
      },
      "state": "failed",
      "status": [
        {
          "aggregate": {
            "name": "aggr1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "error": {
            "code": "852126",
            "message": "string"
          },
          "state": "string"
        }
      ]
```

```
    },
    "interconnect": {
      "adapter": "MVIA-RDMA",
      "state": "string"
    },
    "partners": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "ports": [
      {
        "number": 0,
        "state": "active"
      }
    ],
    "takeover": {
      "failure": {
        "code": 852130,
        "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
      },
      "state": "failed"
    },
    "takeover_check": {
      "reasons": [
        "string"
      ]
    },
    "type": "string"
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "membership": "string",
  "metrocluster": {
```

```
    "ports": [
      {
        "name": "e1b"
      }
    ],
    "type": "string"
  },
  "model": "FAS3070",
  "name": "node-01",
  "nvlog": {
    "backing_type": "string",
    "swap_mode": "string"
  },
  "nvram": {
    "battery_state": "string",
    "id": 0
  },
  "owner": "Example Corp",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "api_service": {
      "port": 0
    },
    "auto_config": {
      "ipv4_subnet": "ipv4_mgmt",
      "ipv6_subnet": "ipv6_mgmt"
    },
    "autoupdate_enabled": true,
    "backup": {
      "state": "string",
      "version": "11.6"
    },
    "dhcp_enabled": true,
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "255.255.0.0",
      "setup_state": "string"
    },
    "last_update_state": "string",
    "link_status": "string",
    "mac_address": "string",
    "primary": {
      "state": "string",
      "version": "11.6"
```

```
      },
      "ssh_info": {
        "allowed_addresses": [
          "10.10.10.7/24"
        ]
      },
      "state": "string",
      "type": "string"
    },
    "snaplock": {
      "compliance_clock_time": "2018-06-04 19:00:00 +0000"
    },
    "state": "string",
    "storage_configuration": "string",
    "system_aggregate": {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "system_id": 92027651,
    "system_machine_type": "7Y56-CTOWW1",
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "vendor_serial_number": 791603000068,
    "vm": {
      "provider_type": "string"
    }
  }
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262245 | The value provided was invalid. |
| 1179795 | A node being added is already in the cluster. |
| 1179813 | Fields set for one node must be set for all nodes. |
| 1179817 | The IP address, subnet mask, and gateway must all be provided for cluster management interface. |
| 1179818 | The IP address and gateway must be of the same family. |
| 1179821 | An IP address and subnet mask conflicts with an existing entry. |
| 1179850 | Internal error. Failed to find the gateway for the default SVM. |
| 1179851 | Failed to set the node management LIF configuration on the node. |
| 1377682 | IPv6 is not enabled in the cluster. |

| Error Code | Description |
| --- | --- |
| 9240591 | The name is not valid. The name is already in use by a cluster node, SVM, or it is the name of the local cluster. |
| 9240594 | An invalid name was provided. |
| 11862025 | The specified node management address and existing node management address belong to two subnets. Create a node management interface after creating the cluster or provide cluster and node management interfaces from the same subnet as the cluster management interface. |
| 131727360 | A node cannot be added to the cluster. This is a generic code, see response message for details. |
| 131727382 | Failed to contact the node using the cluster IP address. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

| Name | Type | Description |
|---|---|---|
| ip | node_setup_ip | The IP configuration for cluster setup. |

ip

IP information

cluster_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
|---|---|---|
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
|---|---|---|
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
|---|---|---|
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |

| Name | Type | Description |
| --- | --- | --- |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |

| Name | Type | Description |
| --- | --- | --- |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
|---|---|---|
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
|---|---|---|
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
|---|---|---|
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
|---|---|---|
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |

hw_assist

The hardware assist information.

| Name | Type | Description |
| --- | --- | --- |
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

| Name | Type | Description |
| --- | --- | --- |
| ip | node_setup_ip | The IP configuration for cluster setup. |

management_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
| --- | --- | --- |
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |
| backup | backup | Provides the properties of the service processor backup partition. |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |

| Name | Type | Description |
| --- | --- | --- |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |

| Name | Type | Description |
|---|---|---|
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
|---|---|---|
| provider_type | string | Cloud provider where the VM is hosted. |

node

Complete node information

| Name | Type | Description |
|---|---|---|
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interface | cluster_interface | The cluster network IP address of the node to be added. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |

| Name | Type | Description |
| --- | --- | --- |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interface | management_interface | The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes. |
| management_interfaces | array[management_interfaces] | |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |

| Name | Type | Description |
| --- | --- | --- |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a node from a cluster

DELETE `/cluster/nodes/{uuid}`

**Introduced In:** 9.7

Deletes a node from the cluster.
Note that before deleting a node from the cluster, you must shut down all of the node's shared resources, such as virtual interfaces to clients. If any of the node's shared resources are still active, the command fails.

#### Optional parameters:

- `force` - Forcibly removes a node that is down and cannot be brought online to remove its shared resources. This flag is set to "false" by default.

#### Related ONTAP commands

- `cluster remove-node`

#### Learn more

- DOC /cluster/nodes

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Set the force flag to "true" to forcibly remove a node that is down and cannot be brought online to remove its shared resources.<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 458755 | Replication service is offline. |
| 458758 | Failed to load job for cluster remove node operation as the job exists. |
| 1179732 | Cannot remove a node in a single-node cluster. |
| 1179735 | Node is not part of a cluster. |
| 1179848 | Operation not supported at the current effective cluster version. |
| 1179849 | Operation failed because the effective cluster version could not be verified. |
| 1182805 | Cannot remove a node from the node network address of the node to be removed. |
| 2293765 | Removing a node only works for nodes not in failover configuration. |

| Error Code | Description |
| --- | --- |
| 2293767 | Node has volumes. Either move or delete them from the node before removing the node. |
| 2293768 | Node is the home node for one or more logical interfaces. |
| 2293769 | Node is the current node for one or more logical interfaces. |
| 2293770 | Node has data logical interfaces configured as target node. |
| 2293787 | Cannot remove node a node with epsilon. |
| 2293788 | Failed to check HA configuration in the cluster. |
| 2293789 | Removing a node only works for nodes not in HA configuration. |
| 2293790 | Failed to check if the node has epsilon. |
| 2293796 | Cluster ring is offline on the node |
| 2293798 | Cannot forcibly remove a node that is online. |
| 2293800 | Node is configured with MetroCluster. |
| 2293801 | Cannot remove node because it has foreign LUN Imports. |
| 2293812 | Node is a member of MetroCluster DR group. |
| 2293813 | Cannot remove a node from the cluster because a controller replacement is in progress. |
| 2293814 | The DELETE operation is not supported until the cluster is upgraded. |
| 2293816 | Cannot remove node because its Storage Encryption devices use authentication keys (AKs) that will not be available to the node after it leaves the cluster. |
| 2293821 | Cannot remove a node that is offline. |
| 2293842 | A node can only be removed if all of its aggregates are home. |
| 2293843 | Cannot determine if all aggregates belonging to the node to be removed are home. |
| 2293844 | Cannot determine if storage failover is enabled. |
| 10551347 | An update is in progress. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve node information**

GET `/cluster/nodes/{uuid}`

**Introduced In:** 9.6

Retrieves information for the node.

#### Related ONTAP commands

- `cluster add-node-status`
- `cluster date show`
- `cluster ha show`
- `network interface show`
- `network port show`
- `storage failover show`
- `system controller show`
- `system node show`
- `system node show-discovered`
- `system service-processor network show`
- `system service-processor show`
- `system service-processor ssh show`
- `system service-processor image show`
- `version`
- `system service-processor api-service show`
- `system service-processor network auto-configuration show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | • format: uuid |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| anti_ransomware_version | string | Anti ransomware version. |

| Name | Type | Description |
| --- | --- | --- |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| hw_assist | hw_assist | The hardware assist information. |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metric | metric | CPU performance for the nodes. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| statistics | statistics | Raw CPU performance for the nodes. |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |

| Name | Type | Description |
| --- | --- | --- |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |
| vm | vm | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anti_ransomware_version": "1.0",
  "cluster_interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "board": "System Board XXVIII",
    "cpu": {
      "count": 20,
      "firmware_release": "string",
      "processor": "string"
    },
    "failed_fan": {
      "count": 1,
      "message": {
        "code": "111411207",
        "message": "There are no failed fans."
      }
    },
    "failed_power_supply": {
      "count": 1,
      "message": {
        "code": "111411208",
        "message": "There are no failed power supplies."
      }
    },
    "flash_cache": [
      {
```

```
        "capacity": 1024000000000,
        "device_id": 0,
        "firmware_file": "X9170_O000Z6300NVM",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": "string",
        "state": "string",
        "type": "string"
      }
    ],
    "memory_size": 1024000000,
    "over_temperature": "string"
  },
  "date": "2019-04-17 11:49:26 -0400",
  "external_cache": {
    "is_enabled": 1,
    "is_hya_enabled": 1,
    "is_rewarm_enabled": 1
  },
  "external_cache_bypass": {
    "enabled": 1,
    "large_read_ops_allow_percent": 100,
    "reset": 1
  },
  "ha": {
    "giveback": {
      "failure": {
        "code": 852126,
        "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
      },
      "state": "failed",
      "status": [
        {
          "aggregate": {
            "_links": {
              "self": {
```

```
                "href": "/api/resourcelink"
              }
            },
            "name": "aggr1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "error": {
            "code": "852126",
            "message": "string"
          },
          "state": "string"
        }
      ]
    },
    "interconnect": {
      "adapter": "MVIA-RDMA",
      "state": "string"
    },
    "partners": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "ports": [
      {
        "number": 0,
        "state": "active"
      }
    ],
    "takeover": {
      "failure": {
        "code": 852130,
        "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
      },
      "state": "failed"
    },
    "takeover_check": {
      "reasons": [
        "string"
```

```
      ]
    },
    "type": "string"
  },
  "hw_assist": {
    "status": {
      "local": {
        "ip": "string",
        "state": "string"
      },
      "partner": {
        "ip": "string",
        "state": "string"
      }
    }
  },
  "location": "rack 2 row 5",
  "management_interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "membership": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "processor_utilization": 13,
    "status": "ok",
    "timestamp": "2017-01-25 11:20:13 +0000",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "metrocluster": {
    "ports": [
```

```
      {
        "name": "e1b"
      }
    ],
    "type": "string"
  },
  "model": "FAS3070",
  "name": "node-01",
  "nvlog": {
    "backing_type": "string",
    "swap_mode": "string"
  },
  "nvram": {
    "battery_state": "string",
    "id": 0
  },
  "owner": "Example Corp",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "api_service": {
      "port": 0
    },
    "auto_config": {
      "ipv4_subnet": "ipv4_mgmt",
      "ipv6_subnet": "ipv6_mgmt"
    },
    "backup": {
      "state": "string",
      "version": "11.6"
    },
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "255.255.0.0",
      "setup_state": "string"
    },
    "ipv6_interface": {
      "address": "fd20:8b1e:b255:5011:10:141:4:97",
      "gateway": "fd20:8b1e:b255:5011:10::1",
      "link_local_ip": "FE80::/10",
      "netmask": 64,
      "router_ip": "2001:0db8:85a3:0000:0000:8a2e:0370:7334",
      "setup_state": "string"
    },
    "last_update_state": "string",
```

```
      "link_status": "string",
      "mac_address": "string",
      "primary": {
        "state": "string",
        "version": "11.6"
      },
      "ssh_info": {
        "allowed_addresses": [
          "10.10.10.7/24"
        ]
      },
      "state": "string",
      "type": "string"
    },
    "snaplock": {
      "compliance_clock_time": "2018-06-04 19:00:00 +0000"
    },
    "state": "string",
    "statistics": {
      "processor_utilization_base": 12345123,
      "processor_utilization_raw": 13,
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "storage_configuration": "string",
    "system_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "system_id": 92027651,
    "system_machine_type": "7Y56-CTOWW1",
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "vendor_serial_number": 791603000068,
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    },
    "vm": {
```

```
    "provider_type": "string"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of CPUs on the node. |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |

| Name | Type | Description |
|---|---|---|
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
|---|---|---|
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
|---|---|---|
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
| --- | --- | --- |
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
| --- | --- | --- |
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
| --- | --- | --- |
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
| --- | --- | --- |
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |
| local | local | |
| partner | partner | |

hw_assist

The hardware assist information.

| Name | Type | Description |
|---|---|---|
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
|---|---|---|
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |

| Name | Type | Description |
| --- | --- | --- |
| autoupdate_enabled | boolean | Indicates whether the service processor can be automatically updated from ONTAP.<br>• Introduced in: 9.10<br>• x-ntap-readModify: true<br>• x-nullable: true |
| backup | backup | Provides the properties of the service processor backup partition. |
| dhcp_enabled | boolean | Set to "true" to use DHCP to configure an IPv4 interface. Do not provide values for address, netmask and gateway when set to "true". |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| ipv6_interface | ipv6_interface | Object to setup an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |

| Name | Type | Description |
| --- | --- | --- |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |

| Name | Type | Description |
|---|---|---|
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
|---|---|---|
| provider_type | string | Cloud provider where the VM is hosted. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update node information**

PATCH `/cluster/nodes/{uuid}`

**Introduced In:** 9.6

Updates the node information or performs shutdown/reboot actions on a node.

#### Related ONTAP commands

- `cluster ha modify`
- `storage failover modify`
- `system node modify`
- `system node reboot`
- `system node power off`
- `system node power on`
- `system service-processor network modify`
- `system service-processor reboot-sp`
- `system service-processor image modify`
- `system service-processor network auto-configuration enable`
- `system service-processor network auto-configuration disable`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | • format: uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| action | string | query | False | The shutdown action shuts the node down and transfers storage control to its HA group if storage failover is enabled. The reboot action reboots the node and transfers storage control to its HA group if storage failover is enabled. The giveback action transfers storage control back to the owner from its HA group.<br>The "power_off" action shuts the node down with the assistance of the service processor. The "power_on" action restores power to the node with the assistance of the service processor.<br><br>• enum: ["shutdown", "reboot", "giveback", "power_off", "power_on", "takeover_check"] |
| shutdown_reboot_reason | string | query | False | Indicates the reason for the reboot or shutdown. This only applies when an action of reboot or shutdown is provided. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| allow_data_outage | boolean | query | False | This only applies when an action of reboot or shutdown is provided. It allows storage failover to be bypassed along with any failures related to maintaining quorum in the cluster.<br>• Default value: |
| service_processor.firmware_image | string | query | False | Service processor image to boot with after a reboot.<br>• Introduced in: 9.10<br>• enum: ["primary", "backup"] |
| service_processor.action | string | query | False | Action used to reboot the service processor (SP).<br>• Introduced in: 9.10<br>• enum: ["reboot"] |
| allow_version_mismatch | boolean | query | False | Applies only when a reboot action is provided. It allows storage failover to be bypassed along with any failures related to software version mismatch.<br>• Introduced in: 9.12<br>• Default value: |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| override_vetoes | boolean | query | False | Applies only when a giveback action is provided. If giveback is vetoed, you must check the EMS messages to determine the cause. Depending on the reason or reasons, you can decide whether you can safely override the vetoes.<br>• Introduced in: 9.13<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |
| system_machine_type | string | OEM system machine type. |

| Name | Type | Description |
| --- | --- | --- |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

**Example request**


```
{
  "anti_ransomware_version": "1.0",
  "cluster_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "board": "System Board XXVIII",
    "cpu": {
      "count": 20,
      "firmware_release": "string",
      "processor": "string"
    },
    "failed_fan": {
      "count": 1,
      "message": {
        "code": "111411207",
        "message": "There are no failed fans."
      }
    },
    "failed_power_supply": {
      "count": 1,
      "message": {
        "code": "111411208",
        "message": "There are no failed power supplies."
      }
    },
    "flash_cache": [
      {
        "capacity": 1024000000000,
        "device_id": 0,
        "firmware_file": "X9170_O000Z6300NVM",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
```

```
      {
        "id": "string",
        "state": "string",
        "type": "string"
      }
    ],
    "memory_size": 1024000000,
    "over_temperature": "string"
  },
  "date": "2019-04-17 11:49:26 -0400",
  "external_cache": {
    "is_enabled": 1,
    "is_hya_enabled": 1,
    "is_rewarm_enabled": 1
  },
  "external_cache_bypass": {
    "enabled": 1,
    "large_read_ops_allow_percent": 100,
    "reset": 1
  },
  "ha": {
    "giveback": {
      "failure": {
        "code": 852126,
        "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
      },
      "state": "failed",
      "status": [
        {
          "aggregate": {
            "name": "aggr1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "error": {
            "code": "852126",
            "message": "string"
          },
          "state": "string"
        }
      ]
    },
    "interconnect": {
      "adapter": "MVIA-RDMA",
      "state": "string"
    },
```

```
    "partners": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "ports": [
      {
        "number": 0,
        "state": "active"
      }
    ],
    "takeover": {
      "failure": {
        "code": 852130,
        "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
      },
      "state": "failed"
    },
    "takeover_check": {
      "reasons": [
        "string"
      ]
    },
    "type": "string"
  },
  "location": "rack 2 row 5",
  "management_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "membership": "string",
  "metrocluster": {
    "ports": [
      {
        "name": "e1b"
      }
    ],
    "type": "string"
  },
  "model": "FAS3070",
  "name": "node-01",
  "nvlog": {
```

```
      "backing_type": "string",
      "swap_mode": "string"
    },
    "nvram": {
      "battery_state": "string",
      "id": 0
    },
    "owner": "Example Corp",
    "serial_number": "4048820-60-9",
    "service_processor": {
      "api_service": {
        "port": 0
      },
      "auto_config": {
        "ipv4_subnet": "ipv4_mgmt",
        "ipv6_subnet": "ipv6_mgmt"
      },
      "backup": {
        "state": "string",
        "version": "11.6"
      },
      "firmware_version": "string",
      "ipv4_interface": {
        "address": "10.10.10.7",
        "gateway": "10.1.1.1",
        "netmask": "255.255.0.0",
        "setup_state": "string"
      },
      "ipv6_interface": {
        "address": "fd20:8b1e:b255:5011:10:141:4:97",
        "gateway": "fd20:8b1e:b255:5011:10::1",
        "link_local_ip": "FE80::/10",
        "netmask": 64,
        "router_ip": "2001:0db8:85a3:0000:0000:8a2e:0370:7334",
        "setup_state": "string"
      },
      "last_update_state": "string",
      "link_status": "string",
      "mac_address": "string",
      "primary": {
        "state": "string",
        "version": "11.6"
      },
      "ssh_info": {
        "allowed_addresses": [
          "10.10.10.7/24"
```

```
      ]
    },
    "state": "string",
    "type": "string"
  },
  "snaplock": {
    "compliance_clock_time": "2018-06-04 19:00:00 +0000"
  },
  "state": "string",
  "storage_configuration": "string",
  "system_aggregate": {
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "system_id": 92027651,
  "system_machine_type": "7Y56-CTOWW1",
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
  "vendor_serial_number": 791603000068,
  "vm": {
    "provider_type": "string"
  }
}
```


**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**


```
{
  "job": {
    "uuid": "string"
  }
}
```


**Response**

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 852046 | HA partner node is not running to do takeover. |
| 852115 | The reboot/shutdown is prevented because LIFs cannot be moved away from the node. |
| 852269 | This operation is not supported on this platform. |
| 3604514 | A reboot or shutdown request is already in progress. |
| 3604515 | Reboot or shutdown of all nodes results in data service failure and client disruption for the entire cluster. Use "allow-data-outage=true" to bypass this check. |
| 9240591 | The name is not valid. The name is already in use by a cluster node, SVM, or it is the name of the local cluster. |
| 9240606 | The reboot/shutdown is prevented due to quorum warnings. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node_setup_ip

The IP configuration for cluster setup.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

cluster_interfaces

Network interface

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

cpu

CPU information.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of CPUs on the node. |

| Name | Type | Description |
| --- | --- | --- |
| firmware_release | string | Firmware release number. Defined by the CPU manufacturer. |
| processor | string | CPU type on the node. |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of chassis fans. |
| message | string | Message describing the current condition of chassis fans. It is only of use when `failed_fan.count` is not zero. |

failed_fan

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Specifies a count of the number of chassis fans that are not operating within the recommended RPM range. |
| message | message | |

message

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code describing the current condition of power supply. |
| message | string | Message describing the state of any power supplies that are currently degraded. It is only of use when `failed_power_supply.count` is not zero. |

failed_power_supply

| Name | Type | Description |
|---|---|---|
| count | integer | Number of failed power supply units. |
| message | message | |

flash_cache

| Name | Type | Description |
|---|---|---|
| capacity | integer | Size in bytes |
| device_id | integer | |
| firmware_file | string | |
| firmware_version | string | |
| hardware_revision | string | |
| model | string | |
| part_number | string | |
| serial_number | string | |
| slot | string | |
| state | string | |

frus

| Name | Type | Description |
|---|---|---|
| id | string | |
| state | string | |
| type | string | |

controller

Controller information

| Name | Type | Description |
|---|---|---|
| board | string | Type of the system board. This is defined by vendor. |
| cpu | cpu | CPU information. |
| failed_fan | failed_fan | |
| failed_power_supply | failed_power_supply | |

| Name | Type | Description |
| --- | --- | --- |
| flash_cache | array[flash_cache] | A list of Flash-Cache devices. Only returned when requested by name. |
| frus | array[frus] | List of FRUs on the node. Only returned when requested by name. |
| memory_size | integer | Memory available on the node, in bytes. |
| over_temperature | string | Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds. |

external_cache

Cache used for buffer management.

| Name | Type | Description |
| --- | --- | --- |
| is_enabled | boolean | Indicates whether the external cache is enabled. |
| is_hya_enabled | boolean | Indicates whether HyA caching is enabled. |
| is_rewarm_enabled | boolean | Indicates whether rewarm is enabled. |
| pcs_size | integer | PCS size in gigabytes. |

external_cache_bypass

External cache bypass management.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether external cache bypass is enabled. |
| large_read_ops_allow_percent | integer | External cache bypass allowed operations percentage for large reads. |

| Name | Type | Description |
| --- | --- | --- |
| reset | boolean | Initiates an external cache bypass threshold reset action. |

failure

Indicates the failure code and message. This property is not supported on the ASA r2 platform.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Message code |
| message | string | Detailed message based on the state. |

aggregate

Aggregate name and UUID.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

error

Indicates the failed aggregate giveback code and message.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Message code. |
| message | string | Detailed message based on the state. |

status

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate name and UUID. |
| error | error | Indicates the failed aggregate giveback code and message. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Giveback state of the aggregate.<br><br>Possible values include no aggregates to giveback(nothing_to_giveback), failed to disable background disk firmware update(BDFU) on source node(failed_bdfu_source),<br><br>giveback delayed as disk firmware update is in progress on source node(delayed_bdfu_source), performing veto checks(running_checks). |

giveback

Represents the state of the node that is giving storage back to its HA partner.

| Name | Type | Description |
| --- | --- | --- |
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |
| status | array[status] | Giveback status of each aggregate. This property is not supported on the ASA r2 platform. |

interconnect

| Name | Type | Description |
| --- | --- | --- |
| adapter | string | HA interconnect device name. |
| state | string | Indicates the HA interconnect status. |

partners

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

ports

| Name | Type | Description |
|---|---|---|
| number | integer | HA port number |
| state | string | HA port state:<br>• *down* - Logical HA link is down.<br>• *initialized* - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port.<br>• *armed* - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port.<br>• *active* - Logical HA link is active.<br>• *reserved* - Logical HA link is active, but the physical link is down. |

takeover

This represents the state of the node that is taking over storage from its HA partner.

| Name | Type | Description |
|---|---|---|
| failure | failure | Indicates the failure code and message. This property is not supported on the ASA r2 platform. |
| state | string | |

takeover_check

The takeover check response.

| Name | Type | Description |
|---|---|---|
| reasons | array[string] | Reasons why the takeover is not possible. |
| takeover_possible | boolean | Indicates whether the takeover is possible. |

ha

| Name | Type | Description |
|---|---|---|
| auto_giveback | boolean | Specifies whether giveback is automatically initiated when the node that owns the storage is ready. |
| enabled | boolean | Specifies whether or not storage failover is enabled. |
| giveback | giveback | Represents the state of the node that is giving storage back to its HA partner. |
| interconnect | interconnect | |
| partners | array[partners] | Nodes in this node's High Availability (HA) group. |
| ports | array[ports] | |
| takeover | takeover | This represents the state of the node that is taking over storage from its HA partner. |
| takeover_check | takeover_check | The takeover check response. |
| type | string | Type of storage. |

local

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |
| state | string | The hardware assist monitor status. |

partner

| Name | Type | Description |
|---|---|---|
| ip | string | The hardware assist IP address. |
| port | integer | The hardware assist port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The hardware assist monitor status. |

status

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether hardware assist is enabled on the node. |

hw_assist

The hardware assist information.

| Name | Type | Description |
| --- | --- | --- |
| status | status | |

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

metric

CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | |

ports

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

metrocluster

Metrocluster

| Name | Type | Description |
| --- | --- | --- |
| custom_vlan_capable | boolean | Indicates whether the MetroCluster over IP platform supports custom VLAN IDs. |
| ports | array[ports] | MetroCluster over IP ports. |
| type | string | The Metrocluster configuration type |

nvlog

Non-volatile write log settings.

| Name | Type | Description |
| --- | --- | --- |
| backing_type | string | Indicates the current NVLog journal backing type. |
| swap_mode | string | Indicates the current NVLog journal swap mode. |

nvram

| Name | Type | Description |
| --- | --- | --- |
| battery_state | string | Specifies status of the NVRAM battery. Possible values:<br>• *battery_ok*<br>• *battery_partially_discharged*<br>• *battery_fully_discharged*<br>• *battery_not_present*<br>• *battery_near_end_of_life*<br>• *battery_at_end_of_life*<br>• *battery_unknown*<br>• *battery_over_charged*<br>• *battery_fully_charged* |
| id | integer | Vendor specific NVRAM ID of the node. |

api_service

Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the SP API service of the SP or BMC is enabled or disabled. When the SP API service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts SP API service access to cluster nodes only. By default, limit_access is set to true. |
| port | integer | Specifies the port number on the SP or BMC used for the SP API service. By default, port 50000 is used. |

auto_config

Provides the properties of the service processor auto configuration.

| Name | Type | Description |
| --- | --- | --- |
| ipv4_subnet | string | Indicates the service processor auto configuration IPv4 subnet name. To enable IPv4 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |
| ipv6_subnet | string | Indicates the service processor auto configuration IPv6 subnet name. To enable IPv6 auto-config give the subnet name, give the value as null or an empty string "" to disable auto-config. |

backup

Provides the properties of the service processor backup partition.

| Name | Type | Description |
|---|---|---|
| is_current | boolean | Indicates whether the service processor is currently booted from the backup partition. |
| state | string | Status of the backup partition. |
| version | string | Firmware version of the backup partition. |

ipv4_interface

Object to set up an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 address |
| enabled | boolean | Indicates whether the IPv4 interfaces is enabled. It expects dhcp_enabled as "true" or values for address, netmask and gateway when set to "true". |
| gateway | string | The IPv4 address of the default router. |
| netmask | string | Input as IPv4 mask (255.255.0.0). Output is always the netmask length. |
| setup_state | string | Indicates the setup state of the interface. |

ipv6_interface

Object to setup an interface along with its default router.

| Name | Type | Description |
|---|---|---|
| address | string | IPv6 address |
| enabled | boolean | Indicates whether the IPv6 interfaces is enabled. It expects values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| gateway | string | The IPv6 address of the default router. |
| is_ipv6_ra_enabled | boolean | Indicates whether IPv6 RA is enabled. |
| link_local_ip | string | Link local IP address. |
| netmask | integer | The IPv6 netmask/prefix length. The default value is 64 with a valid range of 1 to 127. |
| router_ip | string | Router assigned IP address. |
| setup_state | string | Indicates the setup state of the interface. |

primary

Provides the properties of the service processor primary partition.

| Name | Type | Description |
| --- | --- | --- |
| is_current | boolean | Indicates whether the service processor is currently booted from the primary partition. |
| state | string | Status of the primary partition. |
| version | string | Firmware version of the primary partition. |

ssh_info

Service processor SSH allowed IP address configuration applied across the cluster.

| Name | Type | Description |
| --- | --- | --- |
| allowed_addresses | array[string] | Allowed IP addresses |

web_service

Provides the properties of SP or BMC web service.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether the web service of the SP or BMC is enabled or disabled. When the web service is disabled, features such as network-based firmware updates and network-based down node log collection are not available, and the slower serial-interface is used for firmware updates and down node log collection. |
| limit_access | boolean | Restricts web service access to cluster nodes only. By default, limit_access is set to true. |

service_processor

| Name | Type | Description |
| --- | --- | --- |
| api_service | api_service | Provides the properties of the service processor (SP) or baseboard management controller (BMC) API service. |
| auto_config | auto_config | Provides the properties of the service processor auto configuration. |
| autoupdate_enabled | boolean | Indicates whether the service processor can be automatically updated from ONTAP.<br>• Introduced in: 9.10<br>• x-ntap-readModify: true<br>• x-nullable: true |
| backup | backup | Provides the properties of the service processor backup partition. |
| dhcp_enabled | boolean | Set to "true" to use DHCP to configure an IPv4 interface. Do not provide values for address, netmask and gateway when set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| firmware_version | string | The version of firmware installed. |
| ipv4_interface | ipv4_interface | Object to set up an interface along with its default router. |
| ipv6_interface | ipv6_interface | Object to setup an interface along with its default router. |
| is_ip_configured | boolean | Indicates whether the service processor network is configured. |
| last_update_state | string | Provides the "update status" of the last service processor update. |
| link_status | string | |
| mac_address | string | |
| primary | primary | Provides the properties of the service processor primary partition. |
| ssh_info | ssh_info | Service processor SSH allowed IP address configuration applied across the cluster. |
| state | string | |
| type | string | |
| web_service | web_service | Provides the properties of SP or BMC web service. |

snaplock

SnapLock-related properties.

| Name | Type | Description |
| --- | --- | --- |
| compliance_clock_time | string | SnapLock compliance clock time. |

statistics

Raw CPU performance for the nodes.

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_base | integer | Base counter for CPU Utilization. |

| Name | Type | Description |
| --- | --- | --- |
| processor_utilization_raw | integer | Raw CPU utilization for the node. The change in this value over time should be divided by corresponding change in processor_utilization_base, then multiplied by 100 to calculate the percentage CPU utilization for the node. For example: processor_utilization_raw_t2 - processor_utilization_raw_t1) / (processor_utilization_base_t2 - processor_utilization_base_t1 * 100. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

system_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

vm

| Name | Type | Description |
| --- | --- | --- |
| provider_type | string | Cloud provider where the VM is hosted. |

node

Complete node information

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_version | string | Anti ransomware version. |
| cluster_interfaces | array[cluster_interfaces] | |
| controller | controller | Controller information |

| Name | Type | Description |
| --- | --- | --- |
| date | string | The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format.<br>The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.<br>• example: 2019-04-17 11:49:26 -0400<br>• format: date-time<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| external_cache | external_cache | Cache used for buffer management. |
| external_cache_bypass | external_cache_bypass | External cache bypass management. |
| ha | ha | |
| is_spares_low | boolean | Specifies whether or not the node is in spares low condition. |
| location | string | |
| management_interfaces | array[management_interfaces] | |

| Name | Type | Description |
| --- | --- | --- |
| membership | string | Possible values:<br>• *available* - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for *available* to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.<br>• *joining* - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node.<br>• *member* - Nodes that are members have successfully joined the cluster. |
| metrocluster | metrocluster | Metrocluster |
| model | string | |
| name | string | |
| nvlog | nvlog | Non-volatile write log settings. |
| nvram | nvram | |
| owner | string | Owner of the node. |
| serial_number | string | |
| service_processor | service_processor | |
| snaplock | snaplock | SnapLock-related properties. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the node:<br>• *up* - Node is up and operational.<br>• *booting* - Node is booting up.<br>• *down* - Node has stopped or is dumping core.<br>• *taken_over* - Node has been taken over by its HA partner and is not yet waiting for giveback.<br>• *waiting_for_giveback* - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks.<br>• *degraded* - Node has one or more critical services offline.<br>• *unknown* - Node or its HA partner cannot be contacted and there is no information on the node's state. |
| storage_configuration | string | The storage configuration in the system. Possible values:<br>• *mixed_path*<br>• *single_path*<br>• *multi_path*<br>• *tri_path*<br>• *quad_path*<br>• *mixed_path_ha*<br>• *single_path_ha*<br>• *multi_path_ha*<br>• *tri_path_ha*<br>• *quad_path_ha*<br>• *unknown*<br>• *virtual* |
| system_aggregate | system_aggregate | Aggregate |
| system_id | string | |

| Name | Type | Description |
|---|---|---|
| system_machine_type | string | OEM system machine type. |
| uptime | integer | The total time, in seconds, that the node has been up. |
| uuid | string | |
| vendor_serial_number | string | OEM vendor serial number. |
| vm | vm | |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve node historical performance metrics**

GET `/cluster/nodes/{uuid}/metrics`

**Introduced In:** 9.8

Retrieves historical performance metrics for a node.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |
| processor_utilization | integer | query | False | Filter by processor_utilization |
| timestamp | string | query | False | Filter by timestamp |
| uuid | string | path | True | Unique identifier of the node. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "processor_utilization": 13,
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

records

CPU performance for the nodes.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| processor_utilization | integer | Average CPU Utilization for the node |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Cluster NTP

### ONTAP REST API cluster NTP endpoints

#### Overview

ONTAP uses the Network Time Protocol (NTP) for world clock time synchronization of the cluster. Some functional services
require the time to be correct to within one second for all the nodes in the cluster.

The success and speed of this synchronization depends on the number, alignment, and consistent network latency of external
time servers. It is a best practice to configure ONTAP with four independent external time servers.

To aid set up, the Pre-Cluster API of POST /cluster supports a list of NTP time servers
using either the host name, IPv4 address, or IPv6 address.

You can enhance time security by acquiring private keys from external time servers, recording those keys and configuring the entries that match the external time servers to use those keys.

To use NTP symmetric authentication keys (keys), the shared private key must be recorded
first using the /cluster/ntp/keys API associated with the server and enabled to be used.

#### APIs

There are three sets of APIs. The most basic set is part of the /api/cluster APIs, in which a set of NTP servers are provided.
The next two sets are used to manage the NTP servers in more detail and optionally record keys to enable NTP
symmetric authentication.

#### /api/cluster

More details can be found under the documentation for /api/cluster . This API supports a list of NTP servers to start with. It does not take any individual configuration values for the NTP servers themselves.

#### /api/cluster/ntp/servers

You can use this API for a more detailed configuration of NTP servers. You must use this API to set and enable NTP symmetric authentication keys.

#### /api/cluster/ntp/keys

You can use this API to manage shared NTP symmetric keys that are provided by the remote NTP time server by using the key identifier (ID), type of key, and the private shared key.

# Manage cluster NTP keys

## Manage cluster NTP keys

### Overview

You can configure NTP to use shared private keys between ONTAP and trusted external NTP time servers.

You acquire the keys from the external NTP time servers and individual entries created for each unique key. You can use the /cluster/ntp/servers API to associate a key with an external NTP time server used by ONTAP and enable authentication.

## Fields used for adding an NTP shared key

The required fields are:

- `id`
- `digest_type`
- `secret_key`

### Example

```
# Body
create_ntp_key.txt(body):
{
"id": 10,
"digest_type": "sha1",
"value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/keys" -d
"@create_ntp_key.txt"
```

## Retrieve the NTP symmetric authentication keys

GET `/cluster/ntp/keys`

**Introduced In:** 9.7

Retrieves the collection of NTP symmetric authentication keys known by ONTAP that are uniquely indexed by an identifier.

### Related ONTAP commands

- `cluster time-service ntp key show`

**Learn more**

- DOC /cluster/ntp/keys

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| digest_type | string | query | False | Filter by digest_type |
| value | string | query | False | Filter by value |
| id | integer | query | False | Filter by id<br>• Max value: 65535<br>• Min value: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[ntp_key] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "digest_type": "sha1",
      "id": 10,
      "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ntp_key

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an NTP symmetric authentication key entry

POST `/cluster/ntp/keys`

**Introduced In:** 9.7

Creates an NTP symmetric authentication key entry including the type of key using an unused identifier or index number (ID).

#### Required properties

- `id` - Shared symmetric key number (ID).
- `digest_type` - Shared private key cryptographic hash type.
- `value` - Value of shared private key.

#### Related ONTAP commands

- `cluster time-service ntp key create`

#### Learn more

- DOC /cluster/ntp/keys

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br><br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID is included<br>in the NTP cryptographic hash encoded header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

**Example request**

```
{
  "digest_type": "sha1",
  "id": 10,
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}
```

**Response**

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2097187 | Invalid value for an NTP symmetric authentication key. A SHA1 key must be exactly 40 hexadecimal digits. |
| 2097189 | Too many NTP keys have been configured. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ntp_key

| Name | Type | Description |
|---|---|---|
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an NTP key

DELETE `/cluster/ntp/keys/{id}`

**Introduced In:** 9.7

Deletes an NTP key.

#### Related ONTAP commands

- `cluster time-service ntp key delete`

#### Learn more

- DOC /cluster/ntp/keys

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Key identifier |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2097186 | The key cannot be deleted because it is being used by an NTP server. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve NTP symmetric authentication key details

GET `/cluster/ntp/keys/{id}`

**Introduced In:** 9.7

Retrieves the details of a specific NTP symmetric authentication key by numeric identifier or index (ID).

#### Related ONTAP commands

- `cluster time-service ntp key show`

#### Learn more

- DOC /cluster/ntp/keys

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Key identifier |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID is included<br>in the NTP cryptographic hash encoded header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "digest_type": "sha1",
  "id": 10,
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update NTP symmetric authentication key details

PATCH `/cluster/ntp/keys/{id}`

**Introduced In:** 9.7

Updates the details of a specific NTP symmetric authentication key by numeric identifier or index (ID).

#### Required properties

- `digest_type` - Shared private key cryptographic hash type.
- `value` - Value of shared private key.

#### Related ONTAP commands

- `cluster time-service ntp key modify`

#### Learn more

- DOC /cluster/ntp/keys

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Key identifier |

#### Request Body

| Name | Type | Description |
|---|---|---|
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

### Example request

```
{
  "digest_type": "sha1",
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2097187 | An invalid SHA1 key was provided. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ntp_key

| Name | Type | Description |
|---|---|---|
| digest_type | string | The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header. |
| value | string | A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.<br>The current expected length is 40 characters.<br><br>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage cluster NTP servers

### Manage cluster NTP servers

#### Overview

You can use this API to add external NTP servers to a cluster, update the configuration, use NTP keys, and retrieve the
current NTP server configuration.

#### Adding an NTP server to a cluster

To add an NTP server to a cluster, issue a POST /cluster/ntp/servers request.

### Fields used for adding an NTP server

Except for the name of the NTP server (host name or IP address), which is specified by the server, all fields are optional:

- `version`
- `key`

If the key is provided in POST, `authentication_enabled` is set to `true` by default.

#### Examples

### Adding an NTP server

```
# Body
add_ntp_server.txt(body):
{
"server": "time.nist.gov"
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/servers" -d
"@add_ntp_server.txt"
```

---

### Adding an NTP server with an authentication key

```
# Body
add_authenticated_ntp_server.txt(body):
{
"server": "time.nist.gov",
"key": { "id": 10 }
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/servers" -d
"@add_authenticated_ntp_server.txt"
```

---

#### Enabling a previously configured shared key (ID, type, and value) for an NTP server

A combination of key number or identifier (ID), type of key, and shared key value is created with /api/cluster/ntp/keys.
This operation will validate the NTP authentication works.

```
# Body
enable_shared_key.txt(body):
{
"key": { "id": 10 },
"authentication_enabled": true
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/ntp/servers/time.nist.gov" -d
"@enable_shared_key.txt"
```

---

### Retrieve external NTP time servers

GET `/cluster/ntp/servers`

**Introduced In:** 9.7

Retrieves the collection of external NTP time servers ONTAP uses for time adjustment and correction.

#### Related ONTAP commands

- `cluster time-service ntp server show`

#### Learn more

- DOC /cluster/ntp/servers

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| version | string | query | False | Filter by version |
| server | string | query | False | Filter by server |
| authentication_enabled | boolean | query | False | Filter by authentication_enabled |
| key.id | integer | query | False | Filter by key.id<br>• Max value: 65535<br>• Min value: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[ntp_server] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_enabled": 1,
      "key": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 10
      },
      "server": "time.nist.gov",
      "version": "auto"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ntp_key_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server. |

ntp_server

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_enabled | boolean | Set NTP symmetric authentication on (true) or off (false). |
| key | ntp_key_reference | |
| server | string | NTP server host name, IPv4, or IPv6 address. |

| Name | Type | Description |
|---|---|---|
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Validate an external NTP time server**

POST `/cluster/ntp/servers`

**Introduced In:** 9.7

Validates the provided external NTP time server for usage and configures ONTAP so that all nodes in the cluster use it.
The required fields are:

- `server`

#### Default property values

If not specified in POST, the following default property values are assigned:

- `version` - auto
- `key` - not set

If the key is provided in POST, `authentication_enabled` is set to `true` by default.

#### Related ONTAP commands

- `cluster time-service ntp server create`

#### Learn more

- DOC /cluster/ntp/servers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| key | ntp_key_reference | |
| server | string | NTP server host name, IPv4, or IPv6 address. |
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

### Example request

```
{
  "authentication_enabled": 1,
  "key": {
    "id": 10
  },
  "server": "time.nist.gov",
  "version": "auto"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2097163 | NTP server IPv4 address was invalid. |
| 2097164 | NTP server IPv6 address was invalid. |
| 2097165 | Cannot resolve NTP server name. |
| 2097166 | NTP server address query returned no valid IP addresses. |
| 2097167 | Failed to connect to NTP server. |
| 2097168 | Unable to query NTP server. |
| 2097169 | NTP server provided was not synchronized with a clock or another NTP server. |
| 2097174 | NTP server provided had too high of root distance. |
| 2097177 | NTP server provided an invalid stratum. |
| 2097179 | Too many NTP servers have been configured. |
| 2097181 | NTP server address was invalid. It is a special purpose address such as loopback, multicast, or broadcast address. |
| 2097182 | NTP server address was invalid. The address is neither an IPv4 or IPv6. |
| 2097183 | NTP symmetric key authentication cannot be used for a node not in a cluster. |
| 2097185 | NTP key authentication failed for the provided key. |
| 2097193 | An unknown NTP key was provided. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ntp_key_reference

| Name | Type | Description |
|---|---|---|
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server. |

ntp_server

| Name | Type | Description |
|---|---|---|
| key | ntp_key_reference | |
| server | string | NTP server host name, IPv4, or IPv6 address. |
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an external NTP server

DELETE `/cluster/ntp/servers/{server}`

**Introduced In:** 9.7

Deletes an external NTP server used by ONTAP.

#### Related ONTAP commands

- `cluster time-service ntp server delete`

#### Learn more

- DOC /cluster/ntp/servers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| server | string | path | True | Server address or host name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an external NTP server configuration

GET `/cluster/ntp/servers/{server}`

**Introduced In:** 9.7

Retrieves the configuration of an external NTP server used by ONTAP.

#### Related ONTAP commands

- `cluster time-service ntp server show`

#### Learn more

- DOC /cluster/ntp/servers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| server | string | path | True | NTP server host name, IPv4, or IPv6 address. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_enabled | boolean | Set NTP symmetric authentication on (true) or off (false). |
| key | ntp_key_reference | |
| server | string | NTP server host name, IPv4, or IPv6 address. |
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_enabled": 1,
  "key": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 10
  },
  "server": "time.nist.gov",
  "version": "auto"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ntp_key_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an NTP server configuration after validation

PATCH `/cluster/ntp/servers/{server}`

**Introduced In:** 9.7

Updates the configuration of an NTP server used by the ONTAP cluster after validation.
Patchable fields are:

- `version`
- `key.id`
- `authentication_enabled`

If `authentication_enabled` is modified to `false`, the associated NTP key is removed from the server instance.
If `authentication_enabled` is modified to `true`, you must provide an NTP key ID in the PATCH body.

#### Related ONTAP commands

- `cluster time-service ntp server modify`

#### Learn more

- DOC /cluster/ntp/servers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| server | string | path | True | Server address or host name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Set NTP symmetric authentication on (true) or off (false). |
| key | ntp_key_reference | |
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

### Example request

```
{
  "authentication_enabled": 1,
  "key": {
    "id": 10
  },
  "version": "auto"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2097163 | NTP server address was invalid. |
| 2097164 | NTP server address was invalid. |

| Error Code | Description |
| --- | --- |
| 2097165 | Could not resolve NTP server hostname. |
| 2097166 | NTP server address query returned no valid IP addresses. |
| 2097167 | Failed to connect to NTP server. |
| 2097169 | NTP server provided was not synchronized. |
| 2097174 | NTP server provided had too high of root distance. |
| 2097177 | NTP server provided had an invalid stratum. |
| 2097181 | NTP server address was invalid. |
| 2097182 | NTP server address was invalid. |
| 2097183 | NTP symmetric key authentication cannot be used for a node not in a cluster. |
| 2097185 | NTP key authentication failed for the provided key. |
| 2097188 | An invalid key identifier was provided. Identifiers must be in the range from 1 to 65535. |
| 2097193 | An unknown key was provided. |
| 2097194 | The field "authentication_enabled" cannot be false when the field NTP key is given. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ntp_key_reference

| Name | Type | Description |
|---|---|---|
| id | integer | NTP symmetric authentication key identifier or index number (ID). This ID,<br>the type of cryptographic hash, and the cryptographic hash value are all<br>provided by the remote NTP server. |

ntp_server

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Set NTP symmetric authentication on (true) or off (false). |
| key | ntp_key_reference | |
| version | string | NTP protocol version for server. Valid versions are 3, 4, or auto. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster peers

### Manage cluster peers

#### Overview

Cluster peering allows administrators of ONTAP systems to establish relationships between two or more independent clusters. When a relationship exists between two clusters, the clusters can exchange user data and configuration information, and coordinate operations. The /cluster/peers endpoint supports create, get, modify, and delete operations using GET, PATCH, POST and DELETE HTTP requests.

#### Create a cluster peer

You can set up a new cluster peer relationship by issuing a POST request to /cluster/peers. Parameters in the POST body define the settings of the peering relationship. A successful POST request that succeeds in creating a peer returns HTTP status code "201", along with the details of the created peer, such as peer UUID, name, and authentication information. A failed POST request returns an HTTP error code along with a message indicating the reason for the error. This can include malformed requests and invalid operations.

#### Examples of creating cluster peers

### Creating a cluster peer request with an empty request to accept the defaults

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "86de6c46-bdad-11eb-83cd-005056bb267e",
    "name": "Clus_fghf",
    "authentication": {
      "passphrase": "pLznaom1ctesJFq4kt5Qfghf",
      "expiry_time": "2021-05-25T20:04:15-04:00"
    },
    "ip_address": "0.0.0.0",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/86de6c46-bdad-11eb-83cd-005056bb267e"
      }
    }
  }
]
}
```

**Creating a cluster peer request with a system-generated passphrase that will expire on 05/26/2021 at 12:34:56**

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers' -d '{"authentication":
{"expiry_time": "05/26/2021 12:34:56", "generate_passphrase": true}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "14c817c7-bdad-11eb-83cd-005056bb267e",
    "name": "Clus_F6ht",
    "authentication": {
      "passphrase": "dZNOKkpVfntNZHf3MjpNF6ht",
      "expiry_time": "2021-05-26T12:34:56-04:00"
    },
    "ip_address": "0.0.0.0",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/14c817c7-bdad-11eb-83cd-005056bb267e"
      }
    }
  }
]
}
```

**Creating a cluster peer request with a peer address and the generated passphrase is returned in the response**

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers' -d '{"remote":
{"ip_addresses": ["1.2.3.4"]}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "b404cc52-bdae-11eb-812c-005056bb0af1",
    "name": "",
    "authentication": {
      "passphrase": "yDhdOteVGEOhkeXF+DJYwDro",
      "expiry_time": "2021-05-25T20:28:12-04:00"
    },
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b404cc52-bdae-11eb-812c-005056bb0af1"
      }
    }
  }
]
}
```

**Creating a cluster peer request with a peer name and the generated passphrase is returned in the response**

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers' -d '{"name":
"cp_xyz123", "authentication": {"generate_passphrase": true}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "125f8dc6-bdb1-11eb-83cd-005056bb267e",
    "name": "cp_xyz123",
    "authentication": {
      "passphrase": "eeGTerZlh2qSAt2akpYEcM1c",
      "expiry_time": "2021-05-25T20:29:38-04:00"
    },
    "ip_address": "1.2.3.5",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/125f8dc6-bdb1-11eb-83cd-005056bb267e"
      }
    }
  }
]
}
```

**Creating a cluster peer request with a name, a peer address, and a passphrase**

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers' -d '{"name":
"cp_xyz123", "remote": {"ip_addresses": ["1.2.3.4"]}, "authentication":
{"passphrase": "xyz12345"}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "b404cc52-bdae-11eb-812c-005056bb0af1",
    "authentication": {
      "expiry_time": "2021-05-25T20:32:49-04:00"
    },
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b404cc52-bdae-11eb-812c-005056bb0af1"
      }
    }
  }
]
}
```

**Creating a cluster peer request with a proposed encryption protocol**

```
# The API:
/api/cluster/peers

# The call:
curl -X POST 'https://<mgmt-ip>/api/cluster/peers' -d '{"encryption":
{"proposed": "tls-psk"}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "b33a23a6-bdb1-11eb-83cd-005056bb267e",
    "name": "Clus_Pslc",
    "authentication": {
      "passphrase": "Gy8SqsXVhcUkS1AfepH7Pslc",
      "expiry_time": "2021-05-25T20:34:07-04:00"
    },
    "ip_address": "1.2.3.5",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b33a23a6-bdb1-11eb-83cd-005056bb267e"
      }
    }
  }
]
}
```

---

#### Creating local intercluster LIFs

The local cluster must have an intercluster LIF on each node for the correct operation of cluster peering. If no local intercluster LIFs exist, you can optionally specify LIFs to be created for each node in the local cluster. These local interfaces, if specified, are created on each node before proceeding with the creation of the cluster peering relationship. Cluster peering relationships are not established if there is an error preventing the LIFs from being created.
After local interfaces have been created, do not specify them for subsequent cluster peering relationships.

### Local LIF creation fields

- local_network.ip_addresses - List of IP addresses to assign, one per node in the local cluster.
- local_network.netmask - IPv4 mask or subnet mask length.
- local_network.broadcast_domain - Broadcast domain that is in use within the IPspace.
- local_network.gateway - The IPv4 or IPv6 address of the default router.

#### Additional information on network routes

When creating LIFs, the network route discovery mechanism might take additional time (1-5 seconds) to become visible in the network outside of the cluster. This delay in publishing the routes might cause an initial cluster peer "create" request to fail. This error disappears with a retry of the same request.

**This example shows the POST body when creating four intercluster LIFs on a 4-node cluster before creating a cluster peer relationship.**

```
# The API:
/api/cluster/peers

# The call:
cluster_peer_4_node.txt:
{
  "local_network":
  {
      "interfaces": [
          {"ip_address":"1.2.3.4"},
          {"ip_address":"1.2.3.5"},
          {"ip_address":"1.2.3.6"}
      ],
      "netmask": "255.255.0.0",
      "broadcast_domain": "Default",
      "gateway": "1.2.0.1"
  },
  "remote": {"ip_addresses": ["1.2.9.9"]},
  "authentication": {"passphrase": "xyz12345"}
}
curl -X POST "https://<mgmt-ip>/api/cluster/peers" -d
"@cluster_peer_4_node.txt"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "b404cc52-bdae-11eb-812c-005056bb0af1",
    "local_network": {
      "interfaces": [
        {
          "ip_address": "1.2.3.4"
        },
        {
          "ip_address": "1.2.3.5"
        },
        {
```

```
          "ip_address": "1.2.3.6"
        }
      ]
    },
    "authentication": {
      "expiry_time": "2021-05-25T21:28:26-04:00"
    },
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b404cc52-bdae-11eb-812c-005056bb0af1"
      }
    }
  }
]
}
```


#### Examples of retrieving existing cluster peers

You can retrieve peers in a cluster by issuing a GET request to /cluster/peers. It is also possible to retrieve a specific peer when qualified by its UUID to /cluster/peers/{uuid}.
A GET request might have no query parameters or a valid cluster UUID. The former retrieves all records while the latter retrieves the record for the cluster peer with that UUID.

### Retrieving all cluster peer relationships, both established and pending

```
# The API:
/api/cluster/peers

# The call:
curl 'https://<mgmt-ip>/api/cluster/peers'

# The response:
{
"records": [
  {
    "uuid": "a6001076-bdb2-11eb-83cd-005056bb267e",
    "name": "Clus_bH6l",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/a6001076-bdb2-11eb-83cd-005056bb267e"
      },
      "interfaces": {
        "href":
"/api/network/ip/interfaces?services=intercluster_core&ipspace.uuid=0bac5c
ed-a911-11eb-83cd-005056bb267e"
      }
    }
  },
  {
    "uuid": "b404cc52-bdae-11eb-812c-005056bb0af1",
    "name": "remote-cluster",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b404cc52-bdae-11eb-812c-005056bb0af1"
      },
      "interfaces": {
        "href":
"/api/network/ip/interfaces?services=intercluster_core&ipspace.uuid=0bac5c
ed-a911-11eb-83cd-005056bb267e"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/peers"
  }
}
}
```

#### Retrieving all cluster peer relationships which are not in an available state

```
# The API:
/api/cluster/peers

# The call:
curl 'https://<mgmt-ip>/api/cluster/peers?status.state=!available'

# The response:
{
"records": [
  {
    "uuid": "a6001076-bdb2-11eb-83cd-005056bb267e",
    "name": "Clus_bH6l",
    "status": {
      "state": "unidentified"
    },
    "_links": {
      "self": {
        "href": "/api/cluster/peers/a6001076-bdb2-11eb-83cd-005056bb267e"
      },
      "interfaces": {
        "href":
"/api/network/ip/interfaces?services=intercluster_core&ipspace.uuid=0bac5c
ed-a911-11eb-83cd-005056bb267e"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/cluster/peers?status.state=!available"
  }
}
}
```

#### Retrieving information about a single cluster peer relationship

```
# The API:
/api/cluster/peers

# The call:
curl 'https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-812c-
005056bb0af1'
```

```
# The response:
{
"uuid": "b404cc52-bdae-11eb-812c-005056bb0af1",
"name": "remote-cluster",
"version": {
  "full": "NetApp Release 9.10.1: Tue May 25 08:08:44 UTC 2021",
  "generation": 9,
  "major": 10,
  "minor": 1
},
"status": {
  "state": "available",
  "update_time": "2021-05-25T19:38:55-04:00"
},
"ipspace": {
  "uuid": "0bac5ced-a911-11eb-83cd-005056bb267e",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/0bac5ced-a911-11eb-83cd-005056bb267e"
    }
  }
},
"remote": {
  "name": "remote-cluster",
  "serial_number": "1-80-000011",
  "ip_addresses": [
    "1.2.3.4"
  ]
},
"authentication": {
  "in_use": "ok",
  "state": "ok"
},
"encryption": {
  "state": "tls_psk"
},
"_links": {
  "self": {
    "href": "/api/cluster/peers/b404cc52-bdae-11eb-812c-005056bb0af1"
  },
  "interfaces": {
    "href":
"/api/network/ip/interfaces?services=intercluster_core&ipspace.uuid=0bac5c
ed-a911-11eb-83cd-005056bb267e"
```

```
  }
}
}
```

#### Examples of updating an existing cluster peer

You can update a cluster peer relationship by issuing a PATCH request to /cluster/peers/{uuid}. As in the CLI mode, you can toggle the proposed encryption protocol, update the passphrase, or specify a new set of stable addresses. All PATCH requests take the parameters that are to be updated in the request body. If generate_passphrase is "true", the passphrase is returned in the PATCH response.

### Updating the proposed encryption protocol from tls-psk to none

```
# The API:
/api/cluster/peers

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-
812c-005056bb0af1' -d '{"authentication": {"passphrase": "xyz12345",
"in_use": "ok"}, "encryption": {"proposed": "none"}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "authentication": {
      "passphrase": "xyz12345",
      "in_use": "ok"
    },
    "encryption": {
      "proposed": "none"
    }
  }
]
}
```

### Updating the passphrase

```
# The API:
/api/cluster/peers

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-
812c-005056bb0af1' -d '{"authentication": {"passphrase": "xyz12345",
"in_use": "ok"}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "authentication": {
      "passphrase": "xyz12345",
      "in_use": "ok"
    }
  }
]
}
```

**Setting an auto-generated passphrase**

```
# The API:
/api/cluster/peers

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-
812c-005056bb0af1' -d '{"authentication": {"generate_passphrase": true,
"in_use": "ok"}}'

# The response:
{
"num_records": 1,
"records": [
  {
    "authentication": {
      "passphrase": "GpxQquKVLgh2skH+hQ4jRiSf",
      "expiry_time": "2025-03-20T12:11:20-04:00",
      "in_use": "ok"
    }
  }
]
}
```

#### Updating remote IP addresses

```
# The API:
/api/cluster/peers

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-
812c-005056bb0af1' -d '{"remote": {"ip_addresses": ["1.2.3.6"]}}'

# The response:
{}
```

---

#### An example of deleting an existing cluster peer

You can delete a cluster peer using the HTTP DELETE request.

#### Deleting a peer with peer UUID "8becc0d4-c12c-11e8-9ceb-005056bbd143"

```
# The API:
/api/cluster/peers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/peers/b404cc52-bdae-11eb-
812c-005056bb0af1"

# The response:
{}
```

### Retrieve cluster peers

GET `/cluster/peers`

**Introduced In:** 9.6

Retrieves the collection of cluster peers.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cluster_peer] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "interfaces": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication": {
        "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
        "generate_passphrase": true,
        "in_use": "string",
        "passphrase": "string",
        "state": "string"
      },
      "encryption": {
        "proposed": "string",
        "state": "string"
      },
      "initial_allowed_svms": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      ],
      "ip_address": "10.10.10.7",
      "ipspace": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "cluster2",
      "peer_applications": [
        "snapmirror",
        "flexcache"
      ],
      "remote": {
        "ip_addresses": [
          "10.10.10.7"
        ],
        "name": "cluster2",
        "serial_number": "4048820-60-9"
      },
      "status": {
        "state": "available",
        "update_time": "2017-01-25 11:20:13 +0000"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "version": {
        "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
        "generation": 9,
        "major": 4,
        "minor": 0
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| interfaces | href | |
| self | href | |

authentication

| Name | Type | Description |
|---|---|---|
| expiry_time | string | The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour. |
| in_use | string | |
| passphrase | string | A password to authenticate the cluster peer relationship. |
| state | string | |

encryption

| Name | Type | Description |
|---|---|---|
| proposed | string | |
| state | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ipspace

The IPspace of the local intercluster LIFs.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

interfaces

| Name | Type | Description |
|---|---|---|
| ip_address | string | IPv4 or IPv6 address |

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

| Name | Type | Description |
|---|---|---|
| broadcast_domain | string | Broadcast domain that is in use within the IPspace. |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| interfaces | array[interfaces] | |
| netmask | string | IPv4 mask or netmask length. |

remote

| Name | Type | Description |
| --- | --- | --- |
| ip_addresses | array[string] | The IPv4 addresses, IPv6 addresses, or hostnames of the peers. |
| name | string | The name of the remote cluster. |
| serial_number | string | The serial number of the remote cluster. |

status

| Name | Type | Description |
| --- | --- | --- |
| state | string | |
| update_time | string | The last time the state was updated. |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

cluster_peer

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |

| Name | Type | Description |
| --- | --- | --- |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a peering relationship

POST `/cluster/peers`

**Introduced In:** 9.6

Creates a peering relationship and, optionally, the IP interfaces it will use. There are two methods used to create a peering relationship:

- Provide a remote IP address - Used when creating a new cluster peer relationship with a specific remote cluster. This requires at least one remote intercluster IP address from the remote cluster.
- Do not provide a remote IP address - Used when the remote IP address is not provided and when the storage system is ready to accept peering requests from foreign clusters.

#### Required properties

- `remote.ip_addresses` - Addresses of the remote peers. The local peer must be able to reach and connect to these addresses for the request to succeed in creating a peer. Only required when creating a peering relationship by providing a remote IP address.
- Either set `generate_passphrase` to "true" or provide a passphrase in the body of the request. Only one of these options is required.

#### Recommended optional properties

- `name` - Name of the peering relationship or name of the remote peer.
- `passphrase` - User generated passphrase for use in authentication.
- `generate_passphrase` (true/false) - When "true", ONTAP automatically generates a passphrase to authenticate cluster peers.
- `ipspace` - IPspace of the local intercluster LIFs. Assumes Default IPspace if not provided.
- `initial_allowed_svms` - Local SVMs allowed to peer with the peer cluster's SVMs. Can be modified until the remote cluster accepts this cluster peering relationship.
- `local_network` - Fields to create a local intercluster LIF.
- `expiry_time` - Duration in ISO 8601 format for which the user-supplied or auto-generated passphrase is valid. Expiration time must not be greater than seven days into the future. ISO 8601 duration format is "PnDTnHnMnS" or "PnW" where n is a positive integer. The "nD", "nH", "nM" and "nS" fields can be dropped if zero. "P" must always be present and "T" must be present if there are any hours, minutes, or seconds fields.
- `encryption_proposed` (none/tls-psk) - Encryption mechanism of the communication channel between the two peers.

- `peer_applications` - SVM peering applications (SnapMirror, FlexCache or both) for which the SVM peering relationship is set up.

#### Additional information

As with creating a cluster peer through the CLI, the combinations of options must be valid in order for the create operation to succeed. The following list shows the combinations that will succeed and those that will fail:

- A passphrase only (fail)
- A peer IP address (fail)
- A passphrase with an expiration time > 7 days into the future (fail)
- A peer IP address and a passphrase (OK)
- generate_passphrase=true (OK)
- Any proposed encryption protocol (OK)
- An IPspace name or UUID (OK)
- A passphrase, peer IP address, and any proposed encryption protocol (OK)
- A non empty list of initial allowed SVM peer names or UUIDs. (OK)

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |

| Name | Type | Description |
| --- | --- | --- |
| local_network | local_network | Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node. |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |

**Example request**


```
{
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ip_address": "10.10.10.7",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": [
      {
        "ip_address": "10.10.10.7"
      }
    ],
    "netmask": "255.255.0.0"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
  "remote": {
    "ip_addresses": [
      "10.10.10.7"
    ],
    "name": "cluster2",
    "serial_number": "4048820-60-9"
  },
```

```
  "status": {
    "state": "available",
    "update_time": "2017-01-25 11:20:13 +0000"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[cluster_peer] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "authentication": {
        "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
        "in_use": "string",
        "passphrase": "string",
        "state": "string"
      },
      "encryption": {
        "proposed": "string",
        "state": "string"
      },
      "initial_allowed_svms": [
        {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      ],
      "ip_address": "10.10.10.7",
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "local_network": {
        "broadcast_domain": "bd1",
        "gateway": "10.1.1.1",
        "interfaces": [
          {
            "ip_address": "10.10.10.7"
          }
        ],
        "netmask": "255.255.0.0"
      },
      "name": "cluster2",
      "peer_applications": [
        "snapmirror",
        "flexcache"
      ],
      "remote": {
        "ip_addresses": [
          "10.10.10.7"
        ],
```

```
        "name": "cluster2",
        "serial_number": "4048820-60-9"
      },
      "status": {
        "state": "available",
        "update_time": "2017-01-25 11:20:13 +0000"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966366 | The system SVM of the cluster IPspace hosts cluster LIFs only. |
| 4653058 | Cannot peer with self. |
| 4653075 | Cannot peer with two clusters using the same cluster peer relationship name. |
| 4653075 | Cannot peer two clusters with the same name. |
| 4653076 | Cannot peer two clusters with the same UUID. |
| 4653229 | Value for "expiry_time" is obsolete. |
| 4653236 | The specified passphrase is too short. |
| 4653365 | IPspaces are unavailable with cluster peering: {ipspace}. |
| 4653368 | Invalid peer address passed in "ip_addresses": Verify that the peer address is correct, and then try the operation again. |

| Error Code | Description |
| --- | --- |
| 4653419 | An unencrypted peering relationship is not supported because the cluster peering policy requires the use of encryption. |
| 4653708 | The specified addresses must be of the same address family. |
| 4656069 | Specifying a passphrase without remote IP addresses is not supported. |
| 4656070 | The encryption protocol is meaningful only with authenticated cluster peer relationships. |
| 4656071 | Cannot peer with a cluster bearing the same name as the local cluster. |
| 4656072 | The name must conform to the same rules as a cluster name. |
| 4656074 | Cannot check whether all nodes of this cluster support encryption. |
| 4656075 | Cannot specify encryption: this operation requires an ECV of 9.6.0 or later. |
| 4656077 | Specify either remote IP addresses or generate_passphrase. |
| 4656079 | No cluster nodes were found. Check your cluster configuration. |
| 4656081 | Creating an intercluster LIF requires a list of local IP addresses. |
| 4656085 | Cannot create an intercluster LIF with an empty list of local IP addresses. |
| 4656086 | Creating an intercluster LIF requires a broadcast domain that is in use within the IPspace. |
| 4656087 | The number of local intercluster IP addresses must be less than or equal to the number of available nodes. |
| 4656088 | Found no ports matching the IPspace and the broadcast domain. |
| 4656089 | Found no matching entry for IPspace. |
| 4656090 | The given IPspace differs from the IPspace entry found. |
| 4656091 | Creating an intercluster LIF requires a subnet mask or a subnet mask length. |
| 4656094 | Found no ports to bind to intercluster LIFs. Check your network configuration. |

| Error Code | Description |
| --- | --- |
| 4656095 | The address family of the specified peer addresses is not valid in IPspace. Use /api/network/ip/interfaces/ to verify that required LIFs are present and operational on each cluster node. |
| 4656096 | Creating an intercluster LIF requires an IPv4 or IPv6 address of the default router. |
| 4656100 | Cannot use peer address "ip_addresses" because it is already in use on the local cluster. Use /api/network/ip/interfaces to view the address in use. |
| 4656101 | Cannot specify a remote IP that resolves to both IPv4 and IPv6. Specify an IPv4 or IPv6 address. |
| 8847365 | Unknown Host |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

authentication

| Name | Type | Description |
|---|---|---|
| expiry_time | string | The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour. |
| generate_passphrase | boolean | Auto generate a passphrase when true. |
| in_use | string | |
| passphrase | string | A password to authenticate the cluster peer relationship. |
| state | string | |

encryption

| Name | Type | Description |
|---|---|---|
| proposed | string | |
| state | string | |

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ipspace

The IPspace of the local intercluster LIFs.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

interfaces

| Name | Type | Description |
|---|---|---|
| ip_address | string | IPv4 or IPv6 address |

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

| Name | Type | Description |
|---|---|---|
| broadcast_domain | string | Broadcast domain that is in use within the IPspace. |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| interfaces | array[interfaces] | |
| netmask | string | IPv4 mask or netmask length. |

remote

| Name | Type | Description |
|---|---|---|
| ip_addresses | array[string] | The IPv4 addresses, IPv6 addresses, or hostnames of the peers. |
| name | string | The name of the remote cluster. |
| serial_number | string | The serial number of the remote cluster. |

status

| Name | Type | Description |
|---|---|---|
| state | string | |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The last time the state was updated. |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

cluster_peer

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |
| local_network | local_network | Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node. |

| Name | Type | Description |
|---|---|---|
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete a cluster peer**

DELETE `/cluster/peers/{uuid}`

**Introduced In:** 9.6

Deletes a cluster peer.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cluster peer relationship UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4653079 | Unable to delete peer relationship. |
| 4663070 | Unable to delete cluster peer relationship due to an ongoing Vserver migration. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a cluster peer instance

GET `/cluster/peers/{uuid}`

**Introduced In:** 9.6

Retrieves a specific cluster peer instance.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cluster peer relationship UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |
| version | version | This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes. |

**Example response**


```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "generate_passphrase": true,
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ip_address": "10.10.10.7",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
```

```
      "flexcache"
    ],
    "remote": {
      "ip_addresses": [
        "10.10.10.7"
      ],
      "name": "cluster2",
      "serial_number": "4048820-60-9"
    },
    "status": {
      "state": "available",
      "update_time": "2017-01-25 11:20:13 +0000"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    }
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| interfaces | href | |
| self | href | |

authentication

| Name | Type | Description |
|---|---|---|
| expiry_time | string | The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour. |
| in_use | string | |
| passphrase | string | A password to authenticate the cluster peer relationship. |
| state | string | |

encryption

| Name | Type | Description |
|---|---|---|
| proposed | string | |
| state | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ipspace

The IPspace of the local intercluster LIFs.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

interfaces

| Name | Type | Description |
| --- | --- | --- |
| ip_address | string | IPv4 or IPv6 address |

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | string | Broadcast domain that is in use within the IPspace. |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| interfaces | array[interfaces] | |
| netmask | string | IPv4 mask or netmask length. |

remote

| Name | Type | Description |
| --- | --- | --- |
| ip_addresses | array[string] | The IPv4 addresses, IPv6 addresses, or hostnames of the peers. |
| name | string | The name of the remote cluster. |
| serial_number | string | The serial number of the remote cluster. |

status

| Name | Type | Description |
| --- | --- | --- |
| state | string | |
| update_time | string | The last time the state was updated. |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a cluster peer instance

PATCH `/cluster/peers/{uuid}`

**Introduced In:** 9.6

Updates a cluster peer instance.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Cluster peer relationship UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |

**Example request**

```
{
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ip_address": "10.10.10.7",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
  "remote": {
    "ip_addresses": [
      "10.10.10.7"
    ],
    "name": "cluster2",
    "serial_number": "4048820-60-9"
  },
  "status": {
    "state": "available",
    "update_time": "2017-01-25 11:20:13 +0000"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |

**Example response**


```
{
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ip_address": "10.10.10.7",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
  "remote": {
    "ip_addresses": [
      "10.10.10.7"
    ],
    "name": "cluster2",
    "serial_number": "4048820-60-9"
  },
  "status": {
    "state": "available",
    "update_time": "2017-01-25 11:20:13 +0000"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | SVM name does not exist. |
| 4653061 | The specified remote cluster is invalid. |
| 4653215 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653216 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653217 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653218 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653219 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653220 | An introductory RPC to the peer address failed to connect. Verify that the peer address is correct, and then try the operation again. |
| 4653229 | Specified value for "-offer-expiration" is obsolete. |
| 4653236 | The specified passphrase is too short. |
| 4653257 | The vifmgr process is not running. |
| 4653261 | Error finding IPspace. |
| 4653671 | No operational intercluster LIFs of the IPv4 address family is available on this node for the specified IPspace. |
| 4655058 | Expiration time cannot be more than 7 days in the future. |
| 4655061 | SVM does not exist in the IPspace. |
| 4656070 | The encryption protocol is meaningful only with authenticated cluster peer relationships. |
| 4656072 | The name must conform to the same rules as a cluster name. |

| Error Code | Description |
| --- | --- |
| 4656073 | Changing the encryption state requires the refreshing of the authentication passphrase. |
| 4656075 | Cannot specify encryption: this operation requires an ECV of ONTAP 9.6.0 or later. |
| 4656076 | Cluster peer modify was attempted with mismatched IPv4 and IPv6 addresses. |
| 4656080 | Specify either a passphrase or set "generate-passphrase" to true. |
| 4656081 | The remote IP address list is empty. |
| 4656083 | Cannot auto-generate a passphrase when "generate-passphrase" is false. Modifying a passphrase using an auto-generated passphrase requires "generate-passphrase" be true. |
| 4656084 | Passphrase can only be modified with an authenticated cluster peer relationship. |
| 4656092 | Cluster peer modify was attempted with a host name that did not resolve to an IPv4 or IPv6 address. |
| 4656095 | The address family of the specified peer addresses is not valid in this IPspace. Use /api/network/interfaces/ to verify that required LIFs are present and operational on each cluster node. |
| 8847365 | Unknown Host |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

authentication

| Name | Type | Description |
|---|---|---|
| expiry_time | string | The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour. |
| generate_passphrase | boolean | Auto generate a passphrase when true. |
| in_use | string | |
| passphrase | string | A password to authenticate the cluster peer relationship. |
| state | string | |

encryption

| Name | Type | Description |
|---|---|---|
| proposed | string | |
| state | string | |

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ipspace

The IPspace of the local intercluster LIFs.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

interfaces

| Name | Type | Description |
|---|---|---|
| ip_address | string | IPv4 or IPv6 address |

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

| Name | Type | Description |
|---|---|---|
| broadcast_domain | string | Broadcast domain that is in use within the IPspace. |
| gateway | string | The IPv4 or IPv6 address of the default router. |
| interfaces | array[interfaces] | |
| netmask | string | IPv4 mask or netmask length. |

remote

| Name | Type | Description |
|---|---|---|
| ip_addresses | array[string] | The IPv4 addresses, IPv6 addresses, or hostnames of the peers. |
| name | string | The name of the remote cluster. |
| serial_number | string | The serial number of the remote cluster. |

status

| Name | Type | Description |
|---|---|---|
| state | string | |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The last time the state was updated. |

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

| Name | Type | Description |
| --- | --- | --- |
| full | string | The full cluster version string. |
| generation | integer | The generation portion of the version. |
| major | integer | The major portion of the version. |
| minor | integer | The minor portion of the version. |

cluster_peer

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication | authentication | |
| encryption | encryption | |
| initial_allowed_svms | array[initial_allowed_svms] | The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship. |
| ip_address | string | IPv4 or IPv6 address |
| ipspace | ipspace | The IPspace of the local intercluster LIFs. |
| name | string | Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers. |

| Name | Type | Description |
|---|---|---|
| peer_applications | array[string] | Peering applications against which allowed SVMs are configured. |
| remote | remote | |
| status | status | |
| uuid | string | UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster schedules

### Manage cluster schedules

#### Overview

You can use the /cluster/schedules API to view, create, and modify job schedules in a cluster.

#### Retrieving a job schedule

You can retrieve job schedules by issuing a GET request to /cluster/schedules. It is also possible to retrieve a specific schedule when qualified by its UUID to /cluster/schedules/{uuid}. You can apply queries on fields to

retrieve all schedules that match the combined query.

**Example**

```
# The API:
/api/cluster/schedules/

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules?type=interval'

# The response:
{
"records": [
  {
    "uuid": "0941e980-0158-11e9-a82c-005056bb4301",
    "name": "Balanced Placement Model Cache Update",
    "type": "interval",
    "interval": "PT7M30S",
    "_links": {
      "self": {
        "href": "/api/cluster/schedules/0941e980-0158-11e9-a82c-
005056bb4301"
      }
    }
  },
  {
    "uuid": "0944b975-0158-11e9-a82c-005056bb4301",
    "name": "Auto Balance Aggregate Scheduler",
    "type": "interval",
    "interval": "PT1H",
    "_links": {
      "self": {
        "href": "/api/cluster/schedules/0944b975-0158-11e9-a82c-
005056bb4301"
      }
    }
  },
  {
    "uuid": "0c65f1fb-0158-11e9-a82c-005056bb4301",
    "name": "Application Templates ASUP Dump",
    "type": "interval",
    "interval": "P1D",
    "_links": {
      "self": {
        "href": "/api/cluster/schedules/0c65f1fb-0158-11e9-a82c-
005056bb4301"
```

```
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/cluster/schedules?type=interval"
  }
}
}
```

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules/25312bd8-0158-11e9-
a82c-005056bb4301'

# The response:
{
"uuid": "25312bd8-0158-11e9-a82c-005056bb4301",
"name": "monthly",
"cluster": {
  "name": "my_cluster",
  "uuid": "f3f9bbfa-0157-11e9-a82c-005056bb4301"
},
"type": "cron",
"cron": {
  "minutes": [
    20
  ],
  "hours": [
    0
  ],
  "days": [
    1
  ]
},
"_links": {
  "self": {
    "href": "/api/cluster/schedules/25312bd8-0158-11e9-a82c-005056bb4301"
  }
}
}
```

---

#### Creating a job schedule

You can create a job schedule by issuing a POST request to /cluster/schedules to a node in the cluster. For a successful request, the POST request returns a status code of 201.
Job schedules can be of either type "cron" or type "interval". A cron schedule is run at specific minutes within the hour, or hours of the day, days of the week, days of the month, or months of the year. An interval schedule runs repeatedly at fixed intervals.

### Required fields

- name - Name of the job schedule

You are required to provide a "minutes" field for a cron schedule. An "interval" field is required for an interval schedule. Do not provide both a "cron" field and an "interval" field.
The schedule UUID is created by the system.

### Cron schedule fields

- cron.minutes - Minutes within the hour (0 through 59)
- cron.hours - Hours of the day (0 through 23)
- cron.weekdays - Weekdays (0 through 6, where 0 is Sunday and 6 is Saturday.)
- cron.days - Days of the month (1 through 31)
- cron.months - Months of the year (1 through 12)

### Interval schedule field

- interval - Length of time in ISO 8601 duration format.

### Examples

#### Create an interval schedule with a 1-week interval

```
# The API:
/api/cluster/schedules
one_week_interval.txt:
{
  "name": "test_interval_1",
  "interval": "P1W"
}

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d
"@one_week_interval.txt"

# The response of a successful POST is empty.
```

#### Create a cron schedule that runs daily at 12:05

```
# The API:
/api/cluster/schedules
daily_noon_job.txt:
{
  "name": "test_cron_1",
  "cron":
  {
      "minutes": [ 5 ],
      "hours": [ 12 ]
  }
}

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d
"@daily_noon_job.txt"

# The response of a successful POST is empty.
```

### Optional fields

By default, the schedule is owned by the local cluster. In a MetroCluster configuration, you can specify the partner cluster if the local cluster is in the switchover state.

- cluster.name - Name of the cluster owning the schedule.
- cluster.uuid - UUID of the cluster owning the schedule.

### Records field

You can create multiple schedules in one request by providing an array of named records with schedule entries. Each entry must follow the required and optional fields listed above.

---

#### Updating a job schedule

The following fields of an existing schedule can be modified:

- cron.minutes
- cron.hours
- cron.weekdays
- cron.days
- cron.months
- interval
  Note that you cannot modify the name, cluster, and type of schedule. Also, you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule. You can apply queries on fields to modify all schedules that match the combined query.

### Examples

#### Modify an interval schedule with a 2-day and 5-minute interval

```
# The API:
/api/cluster/schedules/{uuid}
every_two_days_five_minutes.txt:
{
  "interval": "P2DT5M"
}

# The call:
curl  -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d
"@every_two_days_five_minutes.txt"

# The response of a successful PATCH is empty.
```

#### Modify a cron schedule to run Mondays at 2

```
# The API:
/api/cluster/schedules/{uuid}
monday_at_two.txt:
{
  "cron":
  {
      "hours": [ 2 ],
      "weekdays": [ 1 ]
  }
}

# The call:
curl  -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d
"@monday_at_two.txt"

# The response of a successful PATCH is empty.
```

---

#### Deleting a job schedule

You can delete job schedules based on their UUID. You can apply queries on fields to delete all schedules that match the combined query.

### Example

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/{uuid}"

# The response of a successful DELETE of one schedule is empty.
```

```
# The API:
/api/cluster/schedules/

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/?name=test*"

# The response of a successful DELETE indicates the number of schedules
affected:
{
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/schedules?name=test*"
  }
}
}
```

---

#### MetroCluster configurations

In a MetroCluster configuration, user-created schedules owned by the local cluster are replicated to the partner cluster. Likewise, user-created schedules owned by the partner cluster are replicated to the local cluster. The owning cluster for a particular schedule is shown in the "cluster.name" and "cluster.uuid" fields.
Normally, only schedules owned by the local cluster can be created, modified, and deleted on the local cluster. However, when a MetroCluster configuration is in switchover, the cluster in switchover state can create, modify, and delete schedules owned by the partner cluster.

### Retrieve schedules

GET `/cluster/schedules`

**Introduced In:** 9.6

Retrieves a schedule.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• maxLength: 256<br>• minLength: 1 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.10 |
| svm.uuid | string | query | False | Filter by svm.uuid<br>• Introduced in: 9.10 |
| type | string | query | False | Filter by type |
| cron.hours | integer | query | False | Filter by cron.hours<br>• Max value: 23<br>• Min value: 0 |
| cron.days | integer | query | False | Filter by cron.days<br>• Max value: 31<br>• Min value: 1 |
| cron.months | integer | query | False | Filter by cron.months<br>• Max value: 12<br>• Min value: 1 |
| cron.minutes | integer | query | False | Filter by cron.minutes<br>• Max value: 59<br>• Min value: 0 |
| cron.weekdays | integer | query | False | Filter by cron.weekdays<br>• Max value: 6<br>• Min value: 0 |
| interval | string | query | False | Filter by interval |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scope | string | query | False | Filter by scope<br>• Introduced in: 9.10 |
| uuid | string | query | False | Filter by uuid |
| cluster.name | string | query | False | Filter by cluster.name |
| cluster.uuid | string | query | False | Filter by cluster.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[schedule] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "cron": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      },
      "interval": "P1DT2H3M4S",
      "name": "string",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 459760 | The schedule specified is not a valid schedule. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

The cluster that owns the schedule. Defaults to the local cluster.

| Name | Type | Description |
|---|---|---|
| name | string | Cluster name |
| uuid | string | Cluster UUID |

cron

Details for schedules of type cron.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | The days of the month the schedule runs. Leave empty for all. |
| hours | array[integer] | The hours of the day the schedule runs. Leave empty for all. |
| minutes | array[integer] | The minutes the schedule runs. Required on POST for a cron schedule. |

| Name | Type | Description |
| --- | --- | --- |
| months | array[integer] | The months of the year the schedule runs. Leave empty for all. |
| weekdays | array[integer] | The weekdays the schedule runs. Leave empty for all. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

schedule

Complete schedule information

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| cluster | cluster | The cluster that owns the schedule. Defaults to the local cluster. |
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| name | string | Schedule name. Required in the URL or POST body. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a schedule

POST `/cluster/schedules`

**Introduced In:** 9.6

Creates a schedule.

#### Required Fields

- name - Name of the job schedule.
  You must provide a minutes field for a cron schedule and an interval field for an interval schedule. Do not provide both a cron field and an interval field.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| cluster | cluster | The cluster that owns the schedule. Defaults to the local cluster. |
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| name | string | Schedule name. Required in the URL or POST body. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

### Example request

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "name": "string",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 458788 | The schedule specified is not a valid schedule. |
| 459760 | The schedule specified is not a valid schedule. |
| 459763 | Schedule cannot be created locally using the remote cluster name as the owner. |
| 459764 | Cannot create a schedule with the same name as an existing schedule from the MetroCluster partner cluster but of a different schedule type. |
| 460783 | As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed. |
| 460784 | An error occurred creating the remote cluster version of this schedule. |
| 2621601 | Cannot create a schedule on a system SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

The cluster that owns the schedule. Defaults to the local cluster.

| Name | Type | Description |
|---|---|---|
| name | string | Cluster name |
| uuid | string | Cluster UUID |

cron

Details for schedules of type cron.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | The days of the month the schedule runs. Leave empty for all. |
| hours | array[integer] | The hours of the day the schedule runs. Leave empty for all. |
| minutes | array[integer] | The minutes the schedule runs. Required on POST for a cron schedule. |
| months | array[integer] | The months of the year the schedule runs. Leave empty for all. |
| weekdays | array[integer] | The weekdays the schedule runs. Leave empty for all. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

schedule

Complete schedule information

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster | The cluster that owns the schedule. Defaults to the local cluster. |
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| name | string | Schedule name. Required in the URL or POST body. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a schedule

DELETE `/cluster/schedules/{uuid}`

**Introduced In:** 9.6

Deletes a schedule.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Schedule UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 459758 | Cannot delete a job schedule that is in use. Remove all references to the schedule, and then try to delete again. |
| 459761 | Schedule cannot be deleted on this cluster because it is replicated from the remote cluster. |
| 459762 | The schedule cannot be deleted because it is a system-level schedule. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a schedule

GET `/cluster/schedules/{uuid}`

**Introduced In:** 9.6

Retrieves a schedule.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Schedule UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster | The cluster that owns the schedule. Defaults to the local cluster. |
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| name | string | Schedule name. Required in the URL or POST body. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "name": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

The cluster that owns the schedule. Defaults to the local cluster.

| Name | Type | Description |
|---|---|---|
| name | string | Cluster name |
| uuid | string | Cluster UUID |

cron

Details for schedules of type cron.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | The days of the month the schedule runs. Leave empty for all. |
| hours | array[integer] | The hours of the day the schedule runs. Leave empty for all. |
| minutes | array[integer] | The minutes the schedule runs. Required on POST for a cron schedule. |
| months | array[integer] | The months of the year the schedule runs. Leave empty for all. |
| weekdays | array[integer] | The weekdays the schedule runs. Leave empty for all. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a schedule

PATCH `/cluster/schedules/{uuid}`

**Introduced In:** 9.6

Updates a schedule. Note that you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Schedule UUID |

**Request Body**

| Name | Type | Description |
|---|---|---|
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

**Example request**

```
{
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "scope": "string",
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 458788 | The schedule specified is not a valid schedule. |
| 459760 | The schedule specified is not a valid schedule. |
| 459761 | Schedule cannot be modified on this cluster because it is replicated from the remote cluster. |
| 460783 | As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed. |
| 461785 | A cron schedule cannot have an interval field. |
| 461786 | An interval schedule cannot have a cron field. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

The cluster that owns the schedule. Defaults to the local cluster.

| Name | Type | Description |
|---|---|---|
| name | string | Cluster name |
| uuid | string | Cluster UUID |

cron

Details for schedules of type cron.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | The days of the month the schedule runs. Leave empty for all. |
| hours | array[integer] | The hours of the day the schedule runs. Leave empty for all. |
| minutes | array[integer] | The minutes the schedule runs. Required on POST for a cron schedule. |
| months | array[integer] | The months of the year the schedule runs. Leave empty for all. |
| weekdays | array[integer] | The weekdays the schedule runs. Leave empty for all. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

schedule

Complete schedule information

| Name | Type | Description |
|---|---|---|
| cron | cron | Details for schedules of type cron. |
| interval | string | An ISO-8601 duration formatted string. |
| scope | string | If the schedule is owned by a data SVM, then the scope is set to svm. Otherwise it will be set to cluster. |
| type | string | Schedule type |
| uuid | string | Job schedule UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Retrieve environment sensors

GET `/cluster/sensors`

**Introduced In:** 9.11

Retrieves Environment Sensors

### Related ONTAP commands

- `system node environment sensors show`

### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| warning_high_threshold | integer | query | False | Filter by warning_high_threshold |
| type | string | query | False | Filter by type |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| discrete_value | string | query | False | Filter by discrete_value |
| critical_high_threshold | integer | query | False | Filter by critical_high_threshold |
| threshold_state | string | query | False | Filter by threshold_state |
| name | string | query | False | Filter by name |
| value | integer | query | False | Filter by value |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| critical_low_threshold | integer | query | False | Filter by critical_low_threshold |
| index | integer | query | False | Filter by index |
| discrete_state | string | query | False | Filter by discrete_state |
| value_units | string | query | False | Filter by value_units |
| warning_low_threshold | integer | query | False | Filter by warning_low_threshold |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[sensors] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "critical_high_threshold": 0,
      "critical_low_threshold": 0,
      "discrete_state": "normal",
      "discrete_value": "ok",
      "index": 0,
      "name": "PVCCSA CPU FD",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "threshold_state": "normal",
      "type": "string",
      "value": 831,
      "value_units": "mV",
      "warning_high_threshold": 0,
      "warning_low_threshold": 0
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

sensors

Environment Sensors

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| critical_high_threshold | integer | Value above which the sensor goes into a critically high state. |
| critical_low_threshold | integer | Value below which the sensor goes into a critically low state. |

| Name | Type | Description |
| --- | --- | --- |
| discrete_state | string | Used to determine whether the sensor is in a normal state or any other failed state based on the value of "discrete_value" field. This field is only applicable for discrete sensors. |
| discrete_value | string | Applies to discrete sensors which do not have an integer value. It can have values like on, off, good, bad, ok. |
| index | integer | Provides the sensor ID. |
| name | string | Name of the sensor. |
| node | node | |
| threshold_state | string | Used to determine whether the sensor is in a normal state or any other failed state. |
| type | string | Used to determine the type of the sensor. |
| value | integer | Provides the sensor reading. |
| value_units | string | Units in which the "value" is measured. Some examples of units are mV, mW*hr, C, RPM. |
| warning_high_threshold | integer | Value above which the sensor goes into a warning high state. |
| warning_low_threshold | integer | Value below which the sensor goes into a warning low state. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage environment sensors

### Manage environment sensors

### Overview

You can use this API to retrieve the details of all platform environment sensors

#### Examples

### Retrieving values of a single sensor

```
# The API:
GET /api/cluster/sensors/{node.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/sensors/{node.uuid}/{index}" -H
"accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"node": {
  "uuid": "19ec0b4a-4a4d-11ec-9036-d039ea4a991a",
  "name": "node1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/19ec0b4a-4a4d-11ec-9036-d039ea4a991a"
    }
  }
},
"index": 1,
"name": "PVCCSA CPU FD",
"type": "voltage",
"value": 831,
"value_units": "mV",
"threshold_state": "normal",
"critical_low_threshold": 297,
"warning_low_threshold": 396,
"warning_high_threshold": 1485,
"critical_high_threshold": 1683,
"_links": {
  "self": {
    "href": "/api/cluster/sensors/19ec0b4a-4a4d-11ec-9036-d039ea4a991a/1"
  }
}
}
```

**Retrieve the environment sensors for a node**

GET `/cluster/sensors/{node.uuid}/{index}`

**Introduced In:** 9.11

Retrieve Environment Sensors

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | string | path | True | Filter by index |
| node.uuid | string | path | True | Filter by node.uuid |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| critical_high_threshold | integer | Value above which the sensor goes into a critically high state. |
| critical_low_threshold | integer | Value below which the sensor goes into a critically low state. |
| discrete_state | string | Used to determine whether the sensor is in a normal state or any other failed state based on the value of "discrete_value" field. This field is only applicable for discrete sensors. |
| discrete_value | string | Applies to discrete sensors which do not have an integer value. It can have values like on, off, good, bad, ok. |
| index | integer | Provides the sensor ID. |
| name | string | Name of the sensor. |
| node | node | |
| threshold_state | string | Used to determine whether the sensor is in a normal state or any other failed state. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Used to determine the type of the sensor. |
| value | integer | Provides the sensor reading. |
| value_units | string | Units in which the "value" is measured. Some examples of units are mV, mW*hr, C, RPM. |
| warning_high_threshold | integer | Value above which the sensor goes into a warning high state. |
| warning_low_threshold | integer | Value below which the sensor goes into a warning low state. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "critical_high_threshold": 0,
  "critical_low_threshold": 0,
  "discrete_state": "normal",
  "discrete_value": "ok",
  "index": 0,
  "name": "PVCCSA CPU FD",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "threshold_state": "normal",
  "type": "string",
  "value": 831,
  "value_units": "mV",
  "warning_high_threshold": 0,
  "warning_low_threshold": 0
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage cluster software

## Manage cluster software

### Overview

You can use the ONTAP cluster software API to retrieve and display relevant information about a software profile, software packages collection, software history collection, and firmware packages collection. This API retrieves the information about all software packages present in the cluster, or a specific software package, or firmware upgrade status.

You can use the POST request to download a software package/firmware from an HTTP or FTP server. The PATCH request provides the option to upgrade the cluster software version. Select the `validate_only` field to validate the package before triggering the update. Set the `version` field to trigger the installation of the package in the cluster. You can pause, resume, or cancel any ongoing software upgrade by selecting `action`. You can use the DELETE request to remove a specific software package present in the cluster.

---

### Examples

#### Retrieving software profile information

The following example shows how to retrieve software and firmware profile information. You can check the validation results after selecting the `validate_only` field. Upgrade progress information is available after an upgrade has started.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software?return_timeout=15" -H
"accept: application/hal+json"

# The response:
{
"validation_results": [
 {
   "update_check": "NFS mounts",
   "status": "warning",
   "issue": {
       "message": "Use NFS hard mounts, if possible.",
     }
   ,
   "action": {
       "message": "Use NFS hard mounts, if possible.",
   }
 }
],
"version": "9.5.0",
```

```
"pending_version": "9.6.0",
"nodes": [
  {
    "node": "Node 1",
    "version": "9.5.0",
    "firmware": {
      "cluster_fw_progress": [
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "abc.zip",
          "update_type": "automatic_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        },
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
```

```
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "xyz.zip",
          "update_type": "manual_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        }
      ],
      "disk": {
        "num_waiting_download": 0,
        "total_completion_estimate": 0,
        "average_duration_per_disk": 120,
        "update_status": "idle"
      },
      "shelf": {
        "update_status": "idle",
        "in_progress_count": 2
      },
      "dqp": {
        "revision": "20200117",
        "version": "3.17",
        "file_name": "qual_devices_v2",
```

```
        "record_count": {
          "drive": 680,
          "alias": 200,
          "device": 29,
          "system": 3
        }
      },
      "sp_bmc": {
        "fw_type": "SP",
        "image": " primary",
        "status": "installed",
        "is_current": true,
        "running_version": "1.2.3.4",
        "autoupdate": false,
        "last_update_status": "passed",
        "start_time": "2018-05-21T09:53:04+05:30",
        "percent_done": 100,
        "end_time": "2018-05-21T09:53:04+05:30",
        "in_progress": false
      }
    }
  }
],
"metrocluster": {
  "progress_summary": {
          "message": "Update paused by user"
   },
  "progress_details": {
          "message": "Installing software image on cluster \"sti70-vsim-
ucs165n_siteA\"."
   },
  "clusters": [
    {
      "name": "sti70-vsim-ucs165n_siteA",
      "uuid": "720f046c-4b13-11e9-9c34-005056ac5626",
      "estimated_duration": 3480,
      "elapsed_duration": 0,
      "state": "waiting"
    },
  ]
},
"state": "in_progress",
"start_time": "2018-05-21T09:53:04+05:30",
"end_time": "2018-05-21T11:53:04+05:30",
"estimated_time": 5220,
"elapsed_time": 2140,
```

```
"update_details": [
  {
    "phase": "Data ONTAP updates",
    "state": "in_progress",
    "estimated_duration": 4620,
    "elapsed_duration": 29,
    "node": {
      "name": "sti70-vsim-ucs165n"
    }
  }
],
"status_details": [
  {
    "name": "do-download-job",
    "state": "completed",
    "issue": {
            "message": "Image update complete",
            "code": 0
     },
    "start_time": "2018-05-21T09:53:04+05:30",
    "end_time": "2018-05-21T11:53:04+05:30",
    "node": {
      "name": "sti70-vsim-ucs165n"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/software/"
  }
}
}
```


### Upgrading the software version

The following example shows how to upgrade cluster software. Set the `version` field to trigger the installation of the package. You can select the `validate_only` field to validate the package before the installation starts. Setting `skip_warning` as `true` ignores the validation warning before the installation starts. Setting the `action` field performs a `pause`, `resume`, or `cancel' operation on an ongoing upgrade. An upgrade can only be resumed if it is in the paused state. Setting` stabilize_minutes `allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes. If` show_validation_details` is set to "true", all validation details will be shown in the output.

You can start the upgrade process at the cluster level. There are no options available to start the upgrade for a specific node or HA pair.

**1. Validating the package and verifying the validation results**

The following example shows how to validate a cluster software package. You must validate the package before the software upgrade. Set the `validate_only` field to `true` to start the validation. You can check for validation results in the GET /cluster/software endpoint.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?validate_only=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0"}'

# The response:
{
"job": {
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
}
```

The call to validate the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-
005056956dfc" -H "accept: application/hal+json"

# The response:
{
"uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
"description": "PATCH /api/cluster/software",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
  }
}
}
```

You can check for validation results in the GET /cluster/software endpoint. The following example shows how to check the validation warnings and errors after setting the `validate_only` field to `true`.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept:
application/hal+json"

# The response:
{
"version": "9.7.0",
"validation_results": [
  {
    "update_check": "High Availability status",
    "status": "error",
    "issue": {
      "message": "Cluster HA is not configured in the cluster. Storage
failover is not enabled on node \"node1\", \"node2\".",
    },
    "action": {
```

```
      "message": "Check cluster HA configuration. Check storage failover
status."
    }
  },
  {
    "update_check": "Manual checks",
    "status": "warning",
    "issue" : {
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or the \"What should I verify before I upgrade
with or without Upgrade Advisor\" section in the \"Upgrade ONTAP\"
documentation for the remaining validation checks that need to be
performed before update. Failing to do so can result in an update failure
or an I/O disruption."
    },
    "action": {
      "message": "Refer to the Upgrade Advisor Plan or the \"What should I
verify before I upgrade with or without Upgrade Advisor\" section in the
\"Upgrade ONTAP\" documentation for the remaining validation checks that
need to be performed before update."
    }
  }
],
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0",
    "firmware": {
      "cluster_fw_progress": [
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "abc.zip",
          "update_type": "automatic_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
```

```
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        },
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "xyz.zip",
          "update_type": "automatic_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
```

```
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        }
      ],
      "disk": {
        "num_waiting_download": 0,
        "total_completion_estimate": 0,
        "average_duration_per_disk": 120,
        "update_status": "idle"
      },
      "shelf": {
        "update_status": "idle",
        "in_progress_count": 2
      },
      "dqp": {
        "revision": "20200117",
        "version": "3.17",
        "file_name": "qual_devices_v2",
        "record_count": {
          "drive": 680,
          "alias": 200,
          "device": 29,
          "system": 3
        }
      },
      "sp_bmc": {
        "fw_type": "SP",
        "image": " primary",
        "status": "installed",
        "is_current": true,
        "running_version": "1.2.3.4",
        "autoupdate": false,
        "last_update_status": "passed",
        "start_time": "2018-05-21T09:53:04+05:30",
        "percent_done": 100,
        "end_time": "2018-05-21T09:53:04+05:30",
        "in_progress": false
      }
    }
  }
],
"state": "failed",
```

```
"elapsed_duration": 56,
"estimated_duration": 600,
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}
```

### 2. Updating the cluster

The following example shows how to initiate a cluster software upgrade. You must validate the package before the software upgrade starts. Set the `skip_warnings` field to `true` to skip validation warnings and start the software package upgrade. You can specify the `stabilize_minutes` value between 1 to 60 minutes. Setting `stabilize_minutes` allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes. If the value of `show_validation_details` is set to "true", then all validation details will be shown in the output. By default, on non-MetroCluster configurations, nodes at the target release will be skipped over. However, this can be disabled by setting `skip_nodes_at_target_version` to `false`.

> it is invalid to set `skip_nodes_at_target_version` to `true` in pause, resume, or cancel actions. `skip_nodes_at_target_version` cannot be set to `true` in MetroCluster configurations.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?skip_warnings=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0"}'

# The response:
{
"job": {
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
}
```

The call to update the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-
005056956dfc" -H "accept: application/hal+json"

# The response:
{
"uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
"description": "PATCH /api/cluster/software",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
  }
}
}
```

You can check the update progress information in the GET /cluster/software endpoint. The following example shows how to check the progress of an update after setting the `skip_warnings` field to `true`. Each node's object also includes information about the firmware update status on the node.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept:
application/hal+json"

# The response:
{
"version": "9.7.0",
"validation_results": [
  {
    "update_check": "Manual checks",
    "status": "warning",
```

```
    "issue" : {
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or the \"What should I verify before I upgrade
with or without Upgrade Advisor\" section in the \"Upgrade ONTAP\"
documentation for the remaining validation checks that need to be
performed before update. Failing to do so can result in an update failure
or an I/O disruption."
    },
    "action": {
      "message": "Refer to the Upgrade Advisor Plan or the \"What should I
verify before I upgrade with or without Upgrade Advisor\" section in the
\"Upgrade ONTAP\" documentation for the remaining validation checks that
need to be performed before update."
    }
  }
],
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0",
    "firmware": {
      "cluster_fw_progress": [
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "abc.zip",
          "update_type": "automated_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 3",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
```

```
              "worker_node": {
                "name": "Node 4",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3498"
              },
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        },
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "xyz.zip",
          "update_type": "automated_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        }
```

```
      ],
      "disk": {
        "num_waiting_download": 0,
        "total_completion_estimate": 0,
        "average_duration_per_disk": 120,
        "update_status": "idle"
      },
      "shelf": {
        "update_status": "idle",
        "in_progress_count": 2
      },
      "dqp": {
        "revision": "20200117",
        "version": "3.17",
        "file_name": "qual_devices_v2",
        "record_count": {
          "drive": 680,
          "alias": 200,
          "device": 29,
          "system": 3
        }
      },
      "sp_bmc": {
        "fw_type": "SP",
        "image": " primary",
        "status": "installed",
        "is_current": true,
        "running_version": "1.2.3.4",
        "autoupdate": false,
        "last_update_status": "passed",
        "start_time": "2018-05-21T09:53:04+05:30",
        "percent_done": 100,
        "end_time": "2018-05-21T09:53:04+05:30",
        "in_progress": false
      }
    }
  }
],
"pending_version": "9.7.0",
"state": "in_progress",
"elapsed_duration": 63,
"estimated_duration": 5220,
"status_details": [
  {
    "name": "do-download-job",
    "status": "running",
```

```
    "issue": {
            "message": "Installing software image.",
            "code": 10551400
     },
    "start_time": "2019-01-14T23:12:14+05:30",
    "end_time": "2019-01-14T23:12:14+05:30",
    "node": {
      "name": "node1"
    }
  },
  {
    "name": "do-download-job",
    "status": "running",
    "issue": {
            "message": "Installing software image.",
            "code": 10551400
    },
    "start_time": "2019-01-14T23:12:14+05:30",
    "end_time": "2019-01-14T23:12:14+05:30",
    "node": {
      "name": "node2"
    }
  }
],
"update_details": [
  {
    "phase": "Data ONTAP updates",
    "status": "in-progress",
    "estimated_duration": 4620,
    "elapsed_duration": 10,
    "node": {
      "name": "node1"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "in-progress",
    "estimated_duration": 4620,
    "elapsed_duration": 10,
    "node": {
      "name": "node2"
    }
  }
],
"_links": {
  "self": {
```

```
    "href": "/api/cluster/software"
  }
}
}
```

In the case of a post update check failure, the details are available under the heading "post_update_checks" in the GET /cluster/software endpoint.
The following example shows how to check the progress of an update after a post update check has failed. Each node's object also includes information about the firmware update status on the node.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept:
application/hal+json"

# The response:
{
"version": "9.7.0",
"validation_results": [
  {
    "update_check": "Manual checks",
    "status": "warning",
    "issue" : {
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or the \"What should I verify before I upgrade
with or without Upgrade Advisor\" section in the \"Upgrade ONTAP\"
documentation for the remaining validation checks that need to be
performed before update. Failing to do so can result in an update failure
or an I/O disruption."
    },
    "action": {
      "message": "Refer to the Upgrade Advisor Plan or the \"What should I
verify before I upgrade with or without Upgrade Advisor\" section in the
\"Upgrade ONTAP\" documentation for the remaining validation checks that
need to be performed before update."
    }
  }
],
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0",
```

```
    "firmware": {
      "cluster_fw_progress": [
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "abc.zip",
          "update_type": "automated_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "working",
              "attempts": 3,
              "message": "<message catalog text>",
              "code": 3
            },
            {
              "worker_node": {
                "name": "Node 2",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "completed",
              "attempts": 3,
              "message": "Error message",
              "code": 0
            }
          ]
        },
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
```

```
        },
        "zip_file_name": "xyz.zip",
        "update_type": "automated_update",
        "update_state": [
          {
            "worker_node": {
              "name": "Node 1",
              "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
            },
            "status": "completed",
            "attempts": 1,
            "message": "Error message",
            "code": 0
          },
          {
            "worker_node": {
              "name": "Node 2",
              "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
            },
            "status": "completed",
            "attempts": "3",
            "message": "Error message",
            "code": 0
          }
        ]
      }
    ],
    "disk": {
      "num_waiting_download": 0,
      "total_completion_estimate": 0,
      "average_duration_per_disk": 120,
      "update_status": "idle"
    },
    "shelf": {
      "update_status": "idle",
      "in_progress_count": 2
    },
    "dqp": {
      "revision": "20200117",
      "version": "3.17",
      "file_name": "qual_devices_v2",
      "record_count": {
        "drive": 680,
        "alias": 200,
        "device": 29,
        "system": 3
```

```
        }
      },
      "sp_bmc": {
        "fw_type": "SP",
        "image": " primary",
        "status": "installed",
        "is_current": "true",
        "running_version": "1.2.3.4",
        "autoupdate": "false",
        "last_update_status": "passed",
        "start_time": "2018-05-21T09:53:04+05:30",
        "percent_done": 100,
        "end_time": "2018-05-21T09:53:04+05:30",
        "in_progress": "yes"
      }
    }
  }
],
"pending_version": "9.7.0",
"state": "in_progress",
"elapsed_duration": 63,
"estimated_duration": 5220,
"status_details": [
  {
    "name": "do-download-job",
    "status": "completed",
    "issue": {
            "message": "Image update complete.",
            "code": 0
     },
    "start_time": "2019-01-14T23:12:14+05:30",
    "end_time": "2019-01-14T23:12:14+05:30",
    "node": {
      "name": "node1"
    }
  },
  {
    "name": "do-download-job",
    "status": "completed",
    "issue": {
            "message": "Image update complete.",
            "code": 0
    },
    "start_time": "2019-01-14T23:12:14+05:30",
    "end_time": "2019-01-14T23:12:14+05:30",
    "node": {
```

```
        "name": "node2"
      }
    }
  ],
  "update_details": [
    {
      "phase": "Data ONTAP updates",
      "status": "completed",
      "estimated_duration": 4620,
      "elapsed_duration": 3120,
      "node": {
        "name": "node1"
      }
    },
    {
      "phase": "Data ONTAP updates",
      "status": "completed",
      "estimated_duration": 4620,
      "elapsed_duration": 3210,
      "node": {
        "name": "node2"
      }
    },
    {
      "phase": "Post-update checks",
      "status": "paused_on_error",
      "estimated_duration": 600,
      "elapsed_duration": 10,
      "node": {
        "name": "node2"
      }
    }
  ],
  "post_update_checks": [
    {
      "update_check": "Aggregate Health Status",
      "status": "error",
      "issue": {
        "message": "Not all aggregates are online"
      },
      "action": {
        "message": "Ensure all aggregates are online."
      }
    },
    {
      "update_check": "HA Health Status",
```

```
    "status": "error",
    "issue": {
      "message": "Storage failover is not enabled on nodes of the
cluster."
    },
    "action": {
      "message": "Ensure storage failover is enabled on all nodes of the
cluster."
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}
```

### 3. Pausing, resuming or canceling an upgrade

The following example shows how to `pause` an ongoing cluster software package upgrade. Set the `action` field to `pause`, `resume`, or `cancel` to pause, resume or cancel the upgrade respectively. Not all update operations support these actions. An update can only be resumed if it is in the paused state.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?action=pause" -H
"accept: application/json" -H "Content-Type: application/hal+json" -d '{
"version": "9.5.0"}'

# The response:
{
"job": {
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
}
```

The call to update the software cluster version and/or firmware version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-
005056956dfc" -H "accept: application/hal+json"

# The response:
{
"uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
"description": "PATCH /api/cluster/software",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
  }
}
}
```

You can check the progress of the upgrade in the GET /cluster/software endpoint. The following example shows how to check the progress of the pause upgrade state after setting the `action` field to `pause`.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept:
application/hal+json"

# The response:
{
"version": "9.7.0",
"validation_results": [
  {
    "update_check": "Manual checks",
    "status": "warning",
    "issue" : {
```

```
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or the \"What should I verify before I upgrade
with or without Upgrade Advisor\" section in the \"Upgrade ONTAP\"
documentation for the remaining validation checks that need to be
performed before update. Failing to do so can result in an update failure
or an I/O disruption."
    },
    "action": {
      "message": "Refer to the Upgrade Advisor Plan or the \"What should I
verify before I upgrade with or without Upgrade Advisor\" section in the
\"Upgrade ONTAP\" documentation for the remaining validation checks that
need to be performed before update."
    }
  }
],
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0",
    "firmware": {
      "cluster_fw_progress": [
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "abc.zip",
          "update_type": "automated_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "status": "complete",
```

```
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        },
        {
          "job": {
            "uuid": "5a21663c-a9a0-11ea-af9a-005056bb44d7",
            "_links": {
              "self": {
                "href": "/api/cluster/jobs/5a21663c-a9a0-11ea-af9a-
005056bb44d7"
              }
            }
          },
          "zip_file_name": "xyz.zip",
          "update_type": "automated_update",
          "update_state": [
            {
              "worker_node": {
                "name": "Node 1",
                "uuid": "fcd40f70-f531-11eb-b235-005056bb3497"
              },
              "status": "failed",
              "attempts": 3,
              "message": "Cannot open the local staging zip file.",
              "code": 2228325
            },
            {
              "status": "complete",
              "attempts": 3,
              "message": "Success",
              "code": 0
            }
          ]
        }
      ],
      "disk": {
        "num_waiting_download": 0,
        "total_completion_estimate": 0,
        "average_duration_per_disk": 120,
        "update_status": "idle"
      },
      "shelf": {
        "update_status": "idle",
```

```
        "in_progress_count": 2
      },
      "dqp": {
        "revision": "20200117",
        "version": "3.17",
        "file_name": "qual_devices_v2",
        "record_count": {
          "drive": 680,
          "alias": 200,
          "device": 29,
          "system": 3
        }
      },
      "sp_bmc": {
        "fw_type": "SP",
        "image": " primary",
        "status": "installed",
        "is_current": true,
        "running_version": "1.2.3.4",
        "autoupdate": false,
        "last_update_status": "passed",
        "start_time": "2018-05-21T09:53:04+05:30",
        "percent_done": 100,
        "end_time": "2018-05-21T09:53:04+05:30",
        "in_progress": false
      }
    }
  }
],
"pending_version": "9.7.0",
"state": "pause_pending",
"elapsed_duration": 103,
"estimated_duration": 5220,
"status_details": [
  {
    "status": "in-progress",
    "issue": {
            "message": "Installing software image.",
            "code": 10551400
     },
    "start_time": "2019-01-08T02:54:36+05:30",
    "node": {
      "name": "node1"
    }
  },
  {
```

```
    "status": "in-progress",
    "issue": {
            "message": "Installing software image.",
            "code": 10551400
     },
    "start_time": "2019-01-08T02:54:36+05:30",
    "node": {
      "name": "node2"
    }
  }
],
"update_details": [
  {
    "phase": "Pre-update checks",
    "status": "completed",
    "estimated_duration": 600,
    "elapsed_duration": 54,
    "node": {
      "name": "node1"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49,
    "node": {
      "name": "node2"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}
```

### Uploading a software/firmware package

The following example shows how to upload a software package.

```
# The API:
/api/cluster/software/upload

# The call:
curl -ku username:password -F "file=@image.tgz" -X POST "https://<mgmt-
ip>/api/cluster/software/upload?return_timeout=0"

# The response:
{
"job": {
"uuid": "12db53fd-8326-11ea-91eb-005056bb16e5",
"_links": {
    "self": {
    "href": "/api/cluster/jobs/12db53fd-8326-11ea-91eb-005056bb16e5"
    }
 }
}
}
```

---

### Retrieve the cluster software profile

GET `/cluster/software`

**Introduced In:** 9.6

Retrieves the software profile of a cluster.

#### Related ONTAP commands

- `cluster image show`
- `cluster image show-update-progress`
- `system node image package show`

#### Learn more

- DOC /cluster/software

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| action | string | User triggered action to apply to the install operation |
| elapsed_duration | integer | Elapsed time during the upgrade or validation operation |
| estimated_duration | integer | Overall estimated time for completion of the upgrade or validation operation. |
| metrocluster | metrocluster | |
| nodes | array[software_node_reference] | List of nodes, active versions, and firmware update progressions. |

| Name | Type | Description |
| --- | --- | --- |
| pending_version | string | Version being installed on the system.<br>• example: ONTAP_X_1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| post_update_checks | array[software_validation_reference] | List of failed post-update checks' warnings, errors, and advice. |
| state | string | Operational state of the upgrade |
| status_details | array[software_status_details_reference] | Display status details. |
| update_details | array[software_update_details_reference] | Display update progress details. |
| validation_results | array[software_validation_reference] | List of validation warnings, errors, and advice. |
| version | string | Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation.<br>• example: ONTAP_X<br>• Introduced in: 9.6<br>• x-nullable: true |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "action": "pause",
  "elapsed_duration": 2140,
  "estimated_duration": 5220,
  "metrocluster": {
    "clusters": [
      {
        "elapsed_duration": 2140,
        "estimated_duration": 3480,
        "name": "cluster_A",
        "state": "in_progress"
      }
    ],
    "progress_details": {
      "message": "Switchover in progress"
    },
    "progress_summary": {
      "message": "MetroCluster updated successfully."
    }
  },
  "nodes": [
    {
      "firmware": {
        "cluster_fw_progress": [
          {
            "job": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "uuid": "string"
            },
            "update_state": [
              {
                "attempts": 3,
                "code": 2228325,
                "message": "Cannot open local staging ZIP file
disk_firmware.zip",
```

```
                "status": "string",
                "worker_node": {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "name": "node1",
                  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                }
              }
            ],
            "update_type": "string",
            "zip_file_name": "disk_firmware.zip"
          }
        ],
        "disk": {
          "average_duration_per_disk": 120,
          "num_waiting_download": 0,
          "total_completion_estimate": 0,
          "update_status": "running"
        },
        "dqp": {
          "file_name": "qual_devices_v3",
          "record_count": {
            "alias": 200,
            "device": 29,
            "drive": 680,
            "system": 3
          },
          "revision": "20200117",
          "version": "3.18"
        },
        "shelf": {
          "in_progress_count": 2,
          "update_status": "running"
        },
        "sp_bmc": {
          "autoupdate": "",
          "end_time": "2020-05-17 20:00:00 +0000",
          "fw_type": "SP",
          "image": "primary",
          "is_current": 1,
          "last_update_state": "passed",
          "percent_done": 100,
          "running_version": "1.2.3.4",
```

```
          "start_time": "2020-05-17 20:00:00 +0000",
          "state": "installed"
        },
        "system_firmware": {
          "boot_media_bios_version": "13.16",
          "boot_media_loader_version": "8.2.0",
          "install_pending": "",
          "running_bios_version": "13.16",
          "running_loader_version": "8.2.0"
        }
      },
      "model": "AFF-A800",
      "name": "node1",
      "software_images": [
        {
          "package": "image.tgz"
        }
      ],
      "version": "ONTAP_X"
    }
  ],
  "pending_version": "ONTAP_X_1",
  "post_update_checks": [
    {
      "action": {
        "message": "Use NFS hard mounts, if possible."
      },
      "issue": {
        "message": "Cluster HA is not configured in the cluster."
      },
      "status": "warning",
      "update_check": "nfs_mounts"
    }
  ],
  "state": "completed",
  "status_details": [
    {
      "action": {
        "message": "string"
      },
      "end_time": "2019-02-02 19:00:00 +0000",
      "issue": {
        "code": 10551399,
        "message": "Image update complete"
      },
      "name": "initialize",
```

```
      "node": {
        "name": "node1"
      },
      "start_time": "2019-02-02 19:00:00 +0000",
      "state": "failed"
    }
  ],
  "update_details": [
    {
      "elapsed_duration": 2100,
      "estimated_duration": 4620,
      "node": {
        "name": "node1"
      },
      "phase": "Post-update checks",
      "state": "failed"
    }
  ],
  "validation_results": [
    {
      "action": {
        "message": "Use NFS hard mounts, if possible."
      },
      "issue": {
        "message": "Cluster HA is not configured in the cluster."
      },
      "status": "warning",
      "update_check": "nfs_mounts"
    }
  ],
  "version": "ONTAP_X"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1048576 | Error initializing software image information. |
| 1048577 | Software image information not initialized. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

software_mcc_reference

| Name | Type | Description |
|---|---|---|
| elapsed_duration | integer | Elapsed duration of update time (in seconds) of MetroCluster. |
| estimated_duration | integer | Estimated duration of update time (in seconds) of MetroCluster. |
| name | string | Name of the site in MetroCluster. |
| state | | Upgrade state of MetroCluster. |

progress_details

| Name | Type | Description |
|---|---|---|
| message | string | MetroCluster update progress details. |

progress_summary

| Name | Type | Description |
|---|---|---|
| message | string | MetroCluster update progress summary. |

metrocluster

| Name | Type | Description |
| --- | --- | --- |
| clusters | array[software_mcc_reference] | List of MetroCluster sites, statuses, and active ONTAP versions.<br>• readOnly: 1<br>• Introduced in: 9.6 |
| progress_details | progress_details | |
| progress_summary | progress_summary | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

worker_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

firmware_update_progress_state

| Name | Type | Description |
| --- | --- | --- |
| attempts | integer | |
| code | integer | Code corresponding to the status message. |
| message | string | Error message returned when a cluster firmware update job fails. |
| status | string | |
| worker_node | worker_node | |

firmware_update_progress

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

| Name | Type | Description |
| --- | --- | --- |
| update_state | array[firmware_update_progress_state] | |
| update_type | string | Specifies the type of update. |
| zip_file_name | string | |

firmware_disk

| Name | Type | Description |
| --- | --- | --- |
| average_duration_per_disk | integer | Average firmware update duration per disk (in seconds). |
| num_waiting_download | integer | The number of disks waiting to download the firmware update. |
| total_completion_estimate | integer | Estimated firmware update duration to completion (in minutes). |
| update_status | | Status of the background disk firmware update. |

firmware_dqp_record_count

| Name | Type | Description |
| --- | --- | --- |
| alias | integer | Alias record count |
| device | integer | Device record count |
| drive | integer | Drive record count |
| system | integer | System record count |

firmware_dqp

| Name | Type | Description |
| --- | --- | --- |
| file_name | string | Firmware file name |
| record_count | firmware_dqp_record_count | |
| revision | string | Firmware revision |
| version | string | Firmware version |

firmware_shelf

| Name | Type | Description |
|---|---|---|
| in_progress_count | integer | |
| update_status | | Status of the shelf firmware update. |

firmware_sp_bmc

| Name | Type | Description |
|---|---|---|
| autoupdate | boolean | |
| end_time | string | |
| fw_type | string | |
| image | string | |
| in_progress | boolean | |
| is_current | boolean | |
| last_update_state | string | |
| percent_done | integer | |
| running_version | string | |
| start_time | string | |
| state | string | |

system_firmware

| Name | Type | Description |
|---|---|---|
| boot_media_bios_version | string | BIOS version that is currently staged on the boot media of the node. |
| boot_media_loader_version | string | LOADER version that is currently staged on the boot media of the node. |
| install_pending | boolean | Flag that indicates if the system firmware install is pending. A true value suggests that a node reboot or cluster upgrade is required for completing installation. |
| running_bios_version | string | BIOS version that is currently installed on the node. |

| Name | Type | Description |
| --- | --- | --- |
| running_loader_version | string | LOADER version that is currently installed on the node. |

firmware

| Name | Type | Description |
| --- | --- | --- |
| cluster_fw_progress | array[firmware_update_progress] | |
| disk | firmware_disk | |
| dqp | firmware_dqp | |
| shelf | firmware_shelf | |
| sp_bmc | firmware_sp_bmc | |
| system_firmware | system_firmware | |

software_images

| Name | Type | Description |
| --- | --- | --- |
| package | string | Package file name. |

software_node_reference

| Name | Type | Description |
| --- | --- | --- |
| firmware | firmware | |
| model | string | Model number of the node. |
| name | string | Name of the node. |
| software_images | array[software_images] | List of software image information. |
| version | string | ONTAP version of the node.<br>• example: ONTAP_X<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

action

| Name | Type | Description |
| --- | --- | --- |
| message | string | Specifies the corrective action to take to resolve an error. |

issue

| Name | Type | Description |
| --- | --- | --- |
| message | string | Details of the error or warning encountered by the update checks. |

software_validation_reference

| Name | Type | Description |
| --- | --- | --- |
| action | action | |
| issue | issue | |
| status | string | Status of the update check. |
| update_check | string | Name of the update check. |

action

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code corresponding the status error |
| message | string | Corrective action to be taken to resolve the status error. |

issue

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code corresponding to update status |
| message | string | Update status details |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of the node to be retrieved for status details. |

software_status_details_reference

| Name | Type | Description |
|---|---|---|
| action | action | |
| end_time | string | End time for each status phase. |
| issue | issue | |
| name | string | Name of the phase to be retrieved for status details. |
| node | node | |
| start_time | string | Start time for each status phase. |
| state | string | Status of the phase |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of the node to be retrieved for update details. |

software_update_details_reference

| Name | Type | Description |
|---|---|---|
| elapsed_duration | integer | Elapsed duration for each update phase |
| estimated_duration | integer | Estimated duration for each update phase |
| node | node | |
| phase | string | Phase details |
| state | string | State of the update phase |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the cluster software version

PATCH `/cluster/software`

**Introduced In:** 9.6

Updates the cluster software version.
Important note:

- Setting 'version' triggers the package installation.
- To validate the package for installation but not perform the installation, use the `validate_only` field on the request.

#### Required properties

- `version` - Software version to be installed on the cluster.

#### Recommended optional parameters

- `validate_only` - Required to validate a software package before an upgrade.
- `skip_warnings` - Used to skip validation warnings when starting a software upgrade.
- `action` - Used to pause, resume, or cancel an ongoing software upgrade.
- `stabilize_minutes` - Specifies a custom value between 1 to 60 minutes that allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes.
- `estimate_only` - Estimates the time duration; does not perform any update.
- `nodes_to_update` - Specifies a subset of the cluster's nodes for update.
- `show_validation_details` - If the value is set to true, then all validation details will be shown in the

output.

- `skip_nodes_at_target_version` - If the value is set to true, then nodes already at the target version will not be upgraded. Defaults to true in non-MetroCluster configurations. Setting this option to false will force all the selected nodes of the cluster to undergo upgrade.

#### Related ONTAP commands

- `cluster image validate`
- `cluster image update`
- `cluster image pause-update`
- `cluster image resume-update`
- `cluster image cancel-update`

#### Learn more

- DOC /cluster/software

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| skip_warnings | boolean | query | False | Ignore warnings and proceed with the install. |
| action | string | query | False | Requests an upgrade to pause, resume, or cancel. Note that not all upgrades support these actions. An upgrade can only be resumed if it is in the paused state. When a request to cancel an upgrade is successful, the upgrade state changes to either `success` or `failure`.<br>• enum: ["pause", "resume", "cancel"] |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| stabilize_minutes | integer | query | False | Sets a custom value between 1 to 60 minutes for the upgrade, allowing each node a specified amount of time to stabilize after a reboot.<br><br>• Introduced in: 9.8 |
| estimate_only | boolean | query | False | Generates an estimate of the time required for the overall update operation for the specified package. No update is performed when this option is used. The default is false.<br><br>• Introduced in: 9.9 |
| force_rolling | boolean | query | False | Forces a rolling upgrade on the cluster. This option is not applicable for a single-node cluster and for a two-node MetroCluster. The default is false.<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nodes_to_update | string | query | False | A comma separated list of node names that are to be updated. The nodes must be a part of an HA pair. The default is all nodes. If the nodes_to_update parameter is an empty string, the upgrade errors out and does not proceed.<br><br>• Introduced in: 9.9 |
| show_validation_details | boolean | query | False | If the value is set to true, then all validation details will be shown in the output.<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| skip_nodes_at_target_version | boolean | query | False | Defaults to true in non-MetroCluster configurations. When set to true, nodes already at the target version will not be upgraded. Setting this field to false will force nodes at the target version to undergo upgrade including migrating LIFs, performing takeover/giveback, and rebooting.<br>It's invalid to set this option to true when the action field is set.<br>It's an invalid option until effective cluster version is 9_16_0.<br><br>• Introduced in: 9.16 |
| pause_after | string | query | False | The pause after specified tasks option. When ANDU is paused user interaction is required to resume the update. The default is none.<br><br>• Introduced in: 9.9<br>• enum: ["none", "takeover_giveback", "all"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| validate_only | boolean | query | False | Validate the operation and its parameters, without actually performing the operation. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| action | string | User triggered action to apply to the install operation |
| elapsed_duration | integer | Elapsed time during the upgrade or validation operation |

| Name | Type | Description |
| --- | --- | --- |
| estimated_duration | integer | Overall estimated time for completion of the upgrade or validation operation. |
| metrocluster | metrocluster | |
| nodes | array[software_node_reference] | List of nodes, active versions, and firmware update progressions. |
| pending_version | string | Version being installed on the system.<br>• example: ONTAP_X_1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| post_update_checks | array[software_validation_reference] | List of failed post-update checks' warnings, errors, and advice. |
| state | string | Operational state of the upgrade |
| status_details | array[software_status_details_reference] | Display status details. |
| update_details | array[software_update_details_reference] | Display update progress details. |
| validation_results | array[software_validation_reference] | List of validation warnings, errors, and advice. |
| version | string | Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation.<br>• example: ONTAP_X<br>• Introduced in: 9.6<br>• x-nullable: true |

**Example request**


```
{
  "action": "pause",
  "elapsed_duration": 2140,
  "estimated_duration": 5220,
  "metrocluster": {
    "clusters": [
      {
        "elapsed_duration": 2140,
        "estimated_duration": 3480,
        "name": "cluster_A",
        "state": "in_progress"
      }
    ],
    "progress_details": {
      "message": "Switchover in progress"
    },
    "progress_summary": {
      "message": "MetroCluster updated successfully."
    }
  },
  "nodes": [
    {
      "model": "AFF-A800",
      "name": "node1",
      "software_images": [
        {
          "package": "image.tgz"
        }
      ],
      "version": "ONTAP_X"
    }
  ],
  "pending_version": "ONTAP_X_1",
  "post_update_checks": [
    {
      "action": {
        "message": "Use NFS hard mounts, if possible."
      },
      "issue": {
        "message": "Cluster HA is not configured in the cluster."
      },
      "status": "warning",
      "update_check": "nfs_mounts"
    }
```

```
  ],
  "state": "completed",
  "status_details": [
    {
      "action": {
        "message": "string"
      },
      "end_time": "2019-02-02 19:00:00 +0000",
      "issue": {
        "code": 10551399,
        "message": "Image update complete"
      },
      "name": "initialize",
      "node": {
        "name": "node1"
      },
      "start_time": "2019-02-02 19:00:00 +0000",
      "state": "failed"
    }
  ],
  "update_details": [
    {
      "elapsed_duration": 2100,
      "estimated_duration": 4620,
      "node": {
        "name": "node1"
      },
      "phase": "Post-update checks",
      "state": "failed"
    }
  ],
  "validation_results": [
    {
      "action": {
        "message": "Use NFS hard mounts, if possible."
      },
      "issue": {
        "message": "Cluster HA is not configured in the cluster."
      },
      "status": "warning",
      "update_check": "nfs_mounts"
    }
  ],
  "version": "ONTAP_X"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2424933 | An internal error occurred while checking MetroCluster information. |
| 2424998 | An internal error occurred while checking MetroCluster configuration. |
| 2425737 | An internal error occurred while checking MetroCluster information. |
| 10551297 | Failed to find an update in progress. |
| 10551298 | Internal error while attempting to pause the upgrade operation. |
| 10551300 | There is no update in progress to be paused. |
| 10551303 | The operation is invalid while the update is not in the paused state. |

| Error Code | Description |
| --- | --- |
| 10551304 | An internal error occurred while canceling the current update operation. |
| 10551306 | A cancel operation is not possible at this time. |
| 10551315 | Failed to find the package in the package repository. |
| 10551319 | An internal error occurred while preparing for the update. |
| 10551320 | Failed to create task on master update manager node. |
| 10551325 | The required package version is not present in RDB. |
| 10551353 | Internal error occurred while running the ANDU pre-validation checks. |
| 10551354 | Invalid operation while the package is being deleted. |
| 10551357 | The operation is invalid while a previous update is still in progress. |
| 10551368 | The list of nodes should contain HA pairs only. |
| 10551391 | An internal error occurred while performing the upgrade operation. |
| 10551395 | ANDU validation check resulted in warnings. |
| 10551396 | ANDU validation checks failed. |
| 10551434 | The options are invalid for a resume operation. |
| 10551524 | The upgrade operation is in the running state, pause the operation before attempting to cancel. |
| 10551558 | Cluster configuration is not supported. |
| 10551564 | Batch mode upgrade is not supported for MetroCluster configurations. |
| 10551565 | Failed to contact the MetroCluster partner cluster, check the MetroCluster configuration then try again. |
| 10551566 | Validation is not complete on the partner MetroCluster cluster. |
| 10551567 | The required package is not available on the partner MetroCluster cluster. |
| 10551572 | Internal error while checking the update status from RDB. Ensure that cluster is healthy and retry the operation. |
| 10551573 | The running update is not in paused state on the MetroCluster partner cluster, so it cannot be resumed. |
| 10551574 | The running update is not in the paused state on the MetroCluster partner cluster, so it cannot be canceled. |
| 10551592 | The nodes selection option is not supported for MetroCluster configurations. |

| Error Code | Description |
| --- | --- |
| 10551599 | Double-hop package is missing. |
| 10551602 | A double-hop upgrade is not supported. |
| 10551606 | An option selected requires a higher effective cluster version. |
| 10551643 | Operation invalid while controller replacement is in progress. |
| 10551656 | Invalid operation while ANDU validation checks are in progress. |
| 10551718 | Invalid operation before running validation checks. |
| 10551719 | ANDU validation results are too old. |
| 10551720 | ANDU validation checks have failed with errors on both clusters in a MetroCluster config. |
| 10551724 | ANDU validation checks were run on a different version from the requested operation. |
| 10551837 | An internal error occurred while generating the AutoSupport prior to starting ANDU. |
| 10551840 | An internal error occurred while generating the AutoSupport prior to starting ANDU. |
| 10551841 | An internal error occurred while generating the AutoSupport prior to starting ANDU. |
| 10551854 | All specified nodes are already running the specified package version. |
| 10551859 | Failed to access the package repository and its backup copy. |
| 10551872 | Package version is not mentioned. |
| 10551897 | All nodes are already running the requested version. |
| 10551898 | Failed to get nodes-to-update list. |
| 10551899 | skip_nodes_at_target_version cannot be specified in a validate_only, pause, resume, or cancel operations. |
| 10551900 | skip_nodes_at_target_version cannot be set to true in MetroCluster configurations. |
| 10551901 | skip_nodes_at_target_version conflicts with another specified field. |
| 10551902 | Failed to determine if the cluster is configured for MetroCluster. |
| 10551908 | skip_nodes_at_target_version is not a valid option until effective cluster version is 9_16_0. |
| 10551910 | The operation is invalid while in the current upgrade state. |

| Error Code | Description |
| --- | --- |
| 13303864 | SnapMirror operation is in progress, retry later. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

software_mcc_reference

| Name | Type | Description |
|---|---|---|
| elapsed_duration | integer | Elapsed duration of update time (in seconds) of MetroCluster. |
| estimated_duration | integer | Estimated duration of update time (in seconds) of MetroCluster. |
| name | string | Name of the site in MetroCluster. |
| state | | Upgrade state of MetroCluster. |

progress_details

| Name | Type | Description |
|---|---|---|
| message | string | MetroCluster update progress details. |

progress_summary

| Name | Type | Description |
|---|---|---|
| message | string | MetroCluster update progress summary. |

metrocluster

| Name | Type | Description |
|---|---|---|
| clusters | array[software_mcc_reference] | List of MetroCluster sites, statuses, and active ONTAP versions.<br>• readOnly: 1<br>• Introduced in: 9.6 |
| progress_details | progress_details | |

| Name | Type | Description |
| --- | --- | --- |
| progress_summary | progress_summary | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

worker_node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

firmware_update_progress_state

| Name | Type | Description |
| --- | --- | --- |
| attempts | integer | |
| code | integer | Code corresponding to the status message. |
| message | string | Error message returned when a cluster firmware update job fails. |
| status | string | |
| worker_node | worker_node | |

firmware_update_progress

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |
| update_state | array[firmware_update_progress_state] | |
| update_type | string | Specifies the type of update. |
| zip_file_name | string | |

firmware_disk

| Name | Type | Description |
|---|---|---|
| average_duration_per_disk | integer | Average firmware update duration per disk (in seconds). |
| num_waiting_download | integer | The number of disks waiting to download the firmware update. |
| total_completion_estimate | integer | Estimated firmware update duration to completion (in minutes). |
| update_status | | Status of the background disk firmware update. |

firmware_dqp_record_count

| Name | Type | Description |
|---|---|---|
| alias | integer | Alias record count |
| device | integer | Device record count |
| drive | integer | Drive record count |
| system | integer | System record count |

firmware_dqp

| Name | Type | Description |
|---|---|---|
| file_name | string | Firmware file name |
| revision | string | Firmware revision |
| version | string | Firmware version |

firmware_shelf

| Name | Type | Description |
|---|---|---|
| in_progress_count | integer | |
| update_status | | Status of the shelf firmware update. |

firmware_sp_bmc

| Name | Type | Description |
| --- | --- | --- |
| autoupdate | boolean | |
| end_time | string | |
| fw_type | string | |
| image | string | |
| in_progress | boolean | |
| is_current | boolean | |
| last_update_state | string | |
| percent_done | integer | |
| running_version | string | |
| start_time | string | |
| state | string | |

system_firmware

| Name | Type | Description |
| --- | --- | --- |
| boot_media_bios_version | string | BIOS version that is currently staged on the boot media of the node. |
| boot_media_loader_version | string | LOADER version that is currently staged on the boot media of the node. |
| install_pending | boolean | Flag that indicates if the system firmware install is pending. A true value suggests that a node reboot or cluster upgrade is required for completing installation. |
| running_bios_version | string | BIOS version that is currently installed on the node. |
| running_loader_version | string | LOADER version that is currently installed on the node. |

firmware

| Name | Type | Description |
| --- | --- | --- |
| cluster_fw_progress | array[firmware_update_progress] | |

software_images

| Name | Type | Description |
|---|---|---|
| package | string | Package file name. |

software_node_reference

| Name | Type | Description |
|---|---|---|
| model | string | Model number of the node. |
| name | string | Name of the node. |
| software_images | array[software_images] | List of software image information. |
| version | string | ONTAP version of the node.<br>• example: ONTAP_X<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

action

| Name | Type | Description |
|---|---|---|
| message | string | Specifies the corrective action to take to resolve an error. |

issue

| Name | Type | Description |
|---|---|---|
| message | string | Details of the error or warning encountered by the update checks. |

software_validation_reference

| Name | Type | Description |
|---|---|---|
| action | action | |
| issue | issue | |
| status | string | Status of the update check. |
| update_check | string | Name of the update check. |

action

| Name | Type | Description |
|---|---|---|
| code | integer | Error code corresponding the status error |
| message | string | Corrective action to be taken to resolve the status error. |

issue

| Name | Type | Description |
|---|---|---|
| code | integer | Error code corresponding to update status |
| message | string | Update status details |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of the node to be retrieved for status details. |

software_status_details_reference

| Name | Type | Description |
|---|---|---|
| action | action | |
| end_time | string | End time for each status phase. |
| issue | issue | |
| name | string | Name of the phase to be retrieved for status details. |
| node | node | |
| start_time | string | Start time for each status phase. |
| state | string | Status of the phase |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the node to be retrieved for update details. |

software_update_details_reference

| Name | Type | Description |
| --- | --- | --- |
| elapsed_duration | integer | Elapsed duration for each update phase |
| estimated_duration | integer | Estimated duration for each update phase |
| node | node | |
| phase | string | Phase details |
| state | string | State of the update phase |

software_reference

| Name | Type | Description |
| --- | --- | --- |
| action | string | User triggered action to apply to the install operation |
| elapsed_duration | integer | Elapsed time during the upgrade or validation operation |
| estimated_duration | integer | Overall estimated time for completion of the upgrade or validation operation. |
| metrocluster | metrocluster | |
| nodes | array[software_node_reference] | List of nodes, active versions, and firmware update progressions. |
| pending_version | string | Version being installed on the system.<br>• example: ONTAP_X_1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| post_update_checks | array[software_validation_reference] | List of failed post-update checks' warnings, errors, and advice. |
| state | string | Operational state of the upgrade |
| status_details | array[software_status_details_reference] | Display status details. |
| update_details | array[software_update_details_reference] | Display update progress details. |
| validation_results | array[software_validation_reference] | List of validation warnings, errors, and advice. |
| version | string | Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation.<br>• example: ONTAP_X<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Upload a software or firmware package located on the local file system

POST `/cluster/software/upload`

**Introduced In:** 9.8

Uploads a software or firmware package located on the local filesystem.

#### Related ONTAP commands

- `cluster image package get`

#### Learn more

- DOC /cluster/software

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| file | file | formData | False | Info specification |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10551359 | Internal error. Failed to process the package after download completed. Try uploading the file again or contact technical support for assistance. |
| 10551373 | Operation took longer than the maximum 1 hour time limit. |
| 10551391 | Internal error. Contact technical support for assistance. |
| 10551392 | Internal error. Contact technical support for assistance. |
| 10551496 | Firmware file already exists. |
| 10551797 | Internal error. Failed to check if file upload is enabled. |
| 10551798 | File upload is disabled. Enable file upload by setting "ApacheUploadEnabled 1" in the web services configuration file or contact technical support for assistance. |
| 10551799 | The requested filename exceeds the max path length. Rename using a shorter name. |
| 10551800 | Internal error. Access permissions restrict file upload. This is likely due to a bad web jail setup. Contact technical support for assistance. |

| Error Code | Description |
| --- | --- |
| 10551801 | Internal error. A read/write error occurred when uploading this file. Contact technical support for assistance |
| 10551802 | An invalid argument was supplied to create a file handle. Try uploading the file again or contact technical support for assistance. |
| 10551803 | An unknown error occurred. Retry file upload operation again or contact technical support for assistance. |
| 10551804 | Internal error. There is not sufficient space in the file upload directory to upload this file. Contact technical support for assistance. |
| 10551805 | Internal error in JAIL setup. Contact technical support for assistance. |
| 10551806 | Internal error. Failed to write to file in the webjail directory. Contact technical support for assistance. |
| 10551807 | The request must only contain a single file. More than one file per request is not supported. |
| 10551808 | The request must be of type multipart/form-data. |
| 10551892 | Package upload failed due to an unhealthy node. Restore the node to a healthy state, then retry the operation. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster software downloads

### Manage cluster software downloads

#### Overview

You can use this API to download a software or firmware package from the server and to check the status of the software or firmware download.

#### Examples

### Downloading the software package

The following example shows how to download the software or firmware package from an HTTP or FTP server. If required, provide the url, username, and password to start the download of the package to the cluster.

```
# The API:
/api/cluster/software/download

# The call:
curl -X POST "https://<mgmt-
ip>/api/cluster/software/download?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/hal+json" -d '{ "url":
"http://server/package", "username": "admin", "password": "*********"}'

# The response:
{
"job": {
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
}
```

---

The call to download the software or firmware package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the job `state` field is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-
005056956dfc" -H "accept: application/hal+json"

# The response:
{
"uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
"description": "POST /api/cluster/software/download",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
  }
}
}
```

#### Checking the progress of the software package being downloaded from an HTTP or FTP server

The following example shows how to retrieve the progress status of the software package that is being downloaded from a HTTP or FTP server.

```
# The API:
/api/cluster/software/download

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/download" -H "accept:
application/hal+json"

# The response:
{
"state": "running",
"message": "Package download in progress",
"code": 10551382,
"_links": {
  "self": {
    "href": "/api/cluster/software/download"
  }
}
}
```

---

### Learn more

- DOC /cluster/software/download

### Retrieve software or firmware download status

GET `/cluster/software/download`

**Introduced In:** 9.7

Retrieves the software or firmware download status.

#### Learn more

- DOC /cluster/software/download

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| code | integer | Code returned corresponds to a download message |
| message | string | Download progress details |
| state | string | Download status of the package |

**Example response**

```
{
  "code": 10551382,
  "message": "Package download in progress",
  "state": "success"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 10551382 | Package download is still running. |
| 10551383 | Operation took longer than the maximum 1 hour time limit. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Download a software or firmware package

POST `/cluster/software/download`

**Introduced In:** 9.6

Downloads a software or firmware package from the server.

#### Required properties

- `url` - URL location of the software package

#### Recommended optional parameters

- `username` - Username of HTTPS/FTP server
- `password` - Password of HTTPS/FTP server

#### Related ONTAP commands

- `cluster image package get`

#### Learn more

- DOC /cluster/software/download

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| code | integer | Code returned corresponds to a download message. |
| message | string | Download progress details. |
| password | string | Password for download |
| state | string | Download status of the package. |

| Name | Type | Description |
|---|---|---|
| url | string | HTTP or FTP URL of the package through a server |
| username | string | Username for download |

#### Example request

```
{
  "code": 10551382,
  "message": "Package download in progress",
  "password": "admin_password",
  "state": "success",
  "url": "http://server/package",
  "username": "admin"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10551327 | Package image with the same name already exists. |
| 10551357 | Cannot perform an update when a previous update is still in progress. |
| 10551359 | Internal error. Failed to process the package after download completed. Try uploading the file again or contact technical support for assistance. |
| 10551381 | Package download failed. |
| 10551382 | Package download is still running. |
| 10551384 | Package download has not started. |
| 10551391 | Internal error. Contact technical support for assistance. |
| 10551392 | Internal error. Contact technical support for assistance. |
| 10551496 | Firmware file already exists. |
| 10551797 | Internal error. Failed to check if file upload is enabled. Contact technical support for assistance. |
| 10551859 | Failed to set primary and secondary nodes to store new image. |
| 10551885 | Package download failed due to invalid url. |
| 10551892 | Package download failed due to an unhealthy node. Restore the node to a healthy state, then retry the operation. |
| 39387137 | Invalid URL syntax was provided. Retry with a valid URL. |
| 39387138 | Unsupported URL scheme provided. Retry with one of FILE://, FTP://, or HTTPS://. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

software_package_download

| Name | Type | Description |
|---|---|---|
| code | integer | Code returned corresponds to a download message. |
| message | string | Download progress details. |
| password | string | Password for download |
| state | string | Download status of the package. |
| url | string | HTTP or FTP URL of the package through a server |
| username | string | Username for download |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster software history

### Manage cluster software history

#### Overview

You can use this API to retrieve the history details for software installation requests.

#### Examples

### Retrieving software installation history information

The following example shows how to:

- Retrieve the software package installation history information.
- Display specific node level software installation history information.
- Provide all the attributes by default in response when the self referential link is not present.

```
# The API:
/api/cluster/software/history

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/history" -H "accept:
application/hal+json"

# The response:
{
"node": {
  "uuid": "58cd3a2b-af63-11e8-8b0d-0050568e7279",
  "name": "sti70-vsim-ucs165n",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/58cd3a2b-af63-11e8-8b0d-0050568e7279"
    }
  }
},
"start_time": "2018-09-03T16:18:46+05:30",
"state": "successful",
"from_version": "9.4.0",
"to_version": "9.5.0",
"end_time": "2018-05-21T10:14:51+05:30"
}
```

---

### Learn more

- DOC /cluster/software/history

### Retrieve software installation requests history details

GET `/cluster/software/history`

**Introduced In:** 9.6

Retrieves the history details for software installation requests.

#### Related ONTAP commands

- `cluster image show-update-history`

#### Learn more

- DOC /cluster/software/history

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.name | string | query | False | Filter by node.name<br>• Introduced in: 9.7 |
| node.uuid | string | query | False | Filter by node.uuid<br>• Introduced in: 9.7 |
| start_time | string | query | False | Filter by start_time<br>• Introduced in: 9.7 |
| from_version | string | query | False | Filter by from_version<br>• Introduced in: 9.7 |
| end_time | string | query | False | Filter by end_time<br>• Introduced in: 9.7 |
| to_version | string | query | False | Filter by to_version<br>• Introduced in: 9.7 |
| state | string | query | False | Filter by state<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[software_history] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "end_time": "2019-02-02 20:00:00 +0000",
      "from_version": "ONTAP_X1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "start_time": "2019-02-02 19:00:00 +0000",
      "state": "successful",
      "to_version": "ONTAP_X2"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

software_history

| Name | Type | Description |
|---|---|---|
| end_time | string | Completion time of this installation request. |
| from_version | string | Previous version of node<br>• example: ONTAP_X1<br>• readOnly: 1<br>• Introduced in: 9.7<br>• x-nullable: true |
| node | node | |
| start_time | string | Start time of this installation request. |
| state | string | Status of this installation request. |

| Name | Type | Description |
| --- | --- | --- |
| to_version | string | Updated version of node<br>• example: ONTAP_X2<br>• readOnly: 1<br>• Introduced in: 9.7<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster software packages

### Manage cluster software packages

#### Overview

You can use this API to retrieve the software packages for a cluster.

#### Examples

### Retrieving cluster software packages information

The following example shows how to retrieve the ONTAP software packages in a cluster.

```
# The API:
/api/cluster/software/packages

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/software/packages?return_records=true&return_timeout=15"
-H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "version": "9.7.0",
    "_links": {
      "self": {
        "href": "/api/cluster/software/packages/9.7.0"
      }
    }
  },
  {
    "version": "9.5.0",
    "_links": {
      "self": {
        "href": "/api/cluster/software/packages/9.5.0"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/software/packages"
  }
}
}
```

#### Retrieves the software package information for a particular version

The following example shows how to retrieve the details of a given cluster software package.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/packages/9.7.0" -H
"accept: application/hal+json"

# The response:
{
"version": "9.7.0",
"create_time": "2018-05-21T10:06:59+05:30",
"_links": {
  "self": {
    "href": "/api/cluster/software/packages/9.7.0"
  }
}
}
```

#### Deleting a cluster software package

The following example shows how to delete a package from the cluster. You need to provide the package version that you want to delete. The software package delete creates a job to perform the delete operation.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/software/packages/9.6.0" -H
"accept: application/hal+json"

# The response:
{
"job": {
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
}
```

The call to delete the package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the job `state` field is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-
005056956dfc" -H "accept: application/hal+json"

# The response:
{
"uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
"description": "DELETE /api/cluster/software/packages/9.6.0",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
  }
}
}
```

---

### Learn more

- DOC /cluster/software/packages

### Retrieve cluster software packages

GET `/cluster/software/packages`

**Introduced In:** 9.6

Retrieves the software packages for a cluster.

#### Related ONTAP commands

- `cluster image package show-repository`

#### Learn more

- DOC /cluster/software/packages

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| create_time | string | query | False | Filter by create_time<br>• Introduced in: 9.7 |
| version | string | query | False | Filter by version<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[software_package] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "create_time": "2019-02-04 19:00:00 +0000",
      "version": "ONTAP_X"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

software_package

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| create_time | string | Indicates when this package was loaded |
| version | string | Version of this package<br>• example: ONTAP_X<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a software package from the cluster

DELETE `/cluster/software/packages/{version}`

**Introduced In:** 9.6

Deletes a software package from the cluster. The delete operation fails if the package is currently installed.

#### Related ONTAP commands

- `cluster image package delete`

#### Learn more

- DOC /cluster/software/packages

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| version | string | path | True | |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10551315 | The repository does not contain any packages. |
| 10551322 | An internal error occurred while retrieving the package cleanup status. |
| 10551323 | An internal error occurred while removing the package information on a node. |
| 10551324 | An internal error occurred while removing the package information on multiple nodes. |
| 10551325 | The required package could not be found in the repository. |
| 10551326 | An error occurred while deleting older package cleanup tasks. |
| 10551346 | A validation is in progress, retry command after it is completed. |
| 10551367 | The repository is busy, retry the operation later. |
| 10551388 | The package delete operation timed out. |

| Error Code | Description |
| --- | --- |
| 10551892 | Package delete failed due to an unhealthy node. Restore the node to a healthy state, then retry the operation. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve the software package information**

GET `/cluster/software/packages/{version}`

**Introduced In:** 9.6

Retrieves the software package information.

#### Related ONTAP commands

- `cluster image package show-repository`

#### Learn more

- DOC /cluster/software/packages

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| version | string | path | True | |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| create_time | string | Indicates when this package was loaded |
| version | string | Version of this package<br>• example: ONTAP_X<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "create_time": "2019-02-04 19:00:00 +0000",
  "version": "ONTAP_X"
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update cluster web configurations

### View and update cluster web configurations

### Overview

You can use this API to update web services configurations and to retrieve current configurations.

#### Retrieving the current web services configuration

The cluster web GET API retrieves the current cluster-wide configuration.

#### Updating the current web services configuration

The cluster web PATCH API updates the current cluster-wide configuration.

Once updated, ONTAP restarts the web services to apply the changes.

When updating the certificate, the certificate UUID of an existing certificate known to ONTAP must be provided. The certificate must be of type "server".

A "client-ca" certificate must be installed on ONTAP to enable "client_enabled".

The following fields can be used to update the cluster-wide configuration:

- enabled
- http_port
- https_port
- http_enabled
- csrf.protection_enabled
- csrf.token.concurrent_limit
- csrf.token.idle_timeout
- csrf.token.max_timeout
- certificate.uuid
- client_enabled
- ocsp_enabled
- hsts.enabled
- hsts.max_age

#### Examples

### Retrieving the cluster-wide web services configuration

```
# API:
GET /api/cluster/web

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/web" -H "accept:
application/hal+json"

# The response:
{
"enabled": true,
"http_port": 80,
"https_port": 443,
"state": "online",
"http_enabled": false,
"csrf": {
  "protection_enabled": true,
  "token": {
    "concurrent_limit": 500,
    "idle_timeout": 900,
    "max_timeout": 650
  }
},
"certificate": {
  "uuid": "a3bb219d-4382-1fe0-9c06-1070568ea23d",
  "name": "cert1",
  "_links": {
    "self": {
      "href": "/api/security/certificates/a3bb219d-4382-1fe0-9c06-
1070568ea23d"
    }
  }
},
"client_enabled": false,
"ocsp_enabled": false,
"hsts": {
  "enabled": true,
  "max_age": 31536000
},
"_links": {
  "self": {
    "href": "/api/cluster/web"
  }
}
}
```

### Updating the cluster-wide web services configuration

```
# The API:
PATCH /api/cluster/web

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/web" -d '{ "https_port": 446,
"csrf": { "token": { "concurrent_limit": 600 } } }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 28 May 2021 09:36:43 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
```

### Retrieve the web services configuration

GET `/cluster/web`

**Introduced In:** 9.10

Retrieves the web services configuration.

#### Related ONTAP commands

- `system services web show`
- `security ssl show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |
| client_enabled | boolean | Indicates whether client authentication is enabled. |
| csrf | csrf | |
| enabled | boolean | Indicates whether remote clients can connect to the web services. |
| hsts | hsts | |
| http_enabled | boolean | Indicates whether HTTP is enabled. |
| http_port | integer | HTTP port for cluster-level web services. |
| https_port | integer | HTTPS port for cluster-level web services. |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |
| per_address_limit | integer | The number of connections that can be processed concurrently from the same remote address. |
| state | string | State of the cluster-level web services. |
| wait_queue_capacity | integer | The maximum size of the wait queue for connections exceeding the per-address-limit. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "csrf": {
    "token": {
      "concurrent_limit": 120,
      "idle_timeout": 200,
      "max_timeout": 1200
    }
  },
  "hsts": {
    "max_age": 31536000
  },
  "per_address_limit": 42,
  "state": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

token

| Name | Type | Description |
|---|---|---|
| concurrent_limit | integer | Maximum number of concurrent CSRF tokens. |
| idle_timeout | integer | Time for which an unused CSRF token is retained, in seconds. |
| max_timeout | integer | Time for which an unused CSRF token, regardless of usage is retained, in seconds. |

csrf

| Name | Type | Description |
|---|---|---|
| protection_enabled | boolean | Indicates whether CSRF protection is enabled. |
| token | token | |

hsts

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether HTTP Strict Transport Security (HSTS) is enabled. |
| max_age | integer | HTTPS Strict Transport Security (HSTS) max-age value in seconds. The maximum time, in seconds, that the browser should remember that a site is only to be accessed using HTTPS. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the web services configuration

PATCH `/cluster/web`

**Introduced In:** 9.10

Updates the web services configuration.

#### Related ONTAP commands

- `system services web modify`
- `security ssl modify`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |
| client_enabled | boolean | Indicates whether client authentication is enabled. |
| csrf | csrf | |
| enabled | boolean | Indicates whether remote clients can connect to the web services. |

| Name | Type | Description |
| --- | --- | --- |
| hsts | hsts | |
| http_enabled | boolean | Indicates whether HTTP is enabled. |
| http_port | integer | HTTP port for cluster-level web services. |
| https_port | integer | HTTPS port for cluster-level web services. |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |
| per_address_limit | integer | The number of connections that can be processed concurrently from the same remote address. |
| state | string | State of the cluster-level web services. |
| wait_queue_capacity | integer | The maximum size of the wait queue for connections exceeding the per-address-limit. |

#### Example request

```
{
  "certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "csrf": {
    "token": {
      "concurrent_limit": 120,
      "idle_timeout": 200,
      "max_timeout": 1200
    }
  },
  "hsts": {
    "max_age": 31536000
  },
  "per_address_limit": 42,
  "state": "string"
}
```

#### Response

```
Status: 200, Ok
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9830406 | Reconfiguration of the web services failed. |
| 9830407 | The web services failed to restart. |
| 9830408 | Reconfiguration and/or restart of the web services failed. |

| Error Code | Description |
| --- | --- |
| 9830442 | Client authentication cannot be enabled without a client ca certificate. |
| 9830463 | The cluster must be fully upgraded before modifying this resource. |
| 9830464 | HTTP cannot be enabled when FIPS is also enabled. |
| 9830483 | The CSRF token timeout is invalid. |
| 9830484 | The maximum concurrent CSRF token count cannot be lower than 100. |
| 9830485 | The CSRF idle timeout cannot be greater than the CSRF absolute timeout. |
| 9830486 | CSRF requires an effective cluster version of 9.7 or later. |
| 9830487 | The HTTP and HTTPS ports must not have the same value. |
| 9830488 | The certificate is not a "server" certificate. |
| 9830489 | The certificate does not exist for the given SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
|---|---|---|
| name | string | Certificate name |
| uuid | string | Certificate UUID |

token

| Name | Type | Description |
|---|---|---|
| concurrent_limit | integer | Maximum number of concurrent CSRF tokens. |
| idle_timeout | integer | Time for which an unused CSRF token is retained, in seconds. |
| max_timeout | integer | Time for which an unused CSRF token, regardless of usage is retained, in seconds. |

csrf

| Name | Type | Description |
|---|---|---|
| protection_enabled | boolean | Indicates whether CSRF protection is enabled. |
| token | token | |

hsts

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether HTTP Strict Transport Security (HSTS) is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| max_age | integer | HTTPS Strict Transport Security (HSTS) max-age value in seconds. The maximum time, in seconds, that the browser should remember that a site is only to be accessed using HTTPS. |

web

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |
| client_enabled | boolean | Indicates whether client authentication is enabled. |
| csrf | csrf | |
| enabled | boolean | Indicates whether remote clients can connect to the web services. |
| hsts | hsts | |
| http_enabled | boolean | Indicates whether HTTP is enabled. |
| http_port | integer | HTTP port for cluster-level web services. |
| https_port | integer | HTTPS port for cluster-level web services. |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |
| per_address_limit | integer | The number of connections that can be processed concurrently from the same remote address. |
| state | string | State of the cluster-level web services. |
| wait_queue_capacity | integer | The maximum size of the wait queue for connections exceeding the per-address-limit. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View tags used for resources

### View tags used for resources

### Overview

You can use this endpoint to list all of the tags that have been used on resources in the API. See the section labelled Tagging Resources for Tracking Purposes to find out more information about how to tag a resource.

#### Examples

The following examples show some ways that this endpoint can be used.

### List all of the used tags

```
# The API:
/api/resource-tags

# The call:
curl -X GET "https://<mgmt-ip>/api/resource-tags"

# The response:
{
"records": [
  {
    "value": "team:accounting",
    "num_resources": 2,
    "_links": {
      "self": {
        "href": "/api/resource-tags/team%3Aaccounting"
      }
    }
  },
  {
    "value": "environment:test",
    "num_resources": 5,
    "_links": {
      "self": {
        "href": "/api/resource-tags/environment%3Atest"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/resource-tags"
  }
}
}
```

---

**Find tags that are being used by at least 3 resources**

```
# The API:
/api/resource-tags

# The call:
curl -X GET "https://<mgmt-ip>/api/resource-tags?num_resources=>=3"

# The response:
{
"records": [
  {
    "value": "environment:test",
    "num_resources": 5,
    "_links": {
      "self": {
        "href": "/api/resource-tags/environment%3Atest"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/resource-tags?num_resources=>=3"
  }
}
}
```

**Retrieve the tags currently used for resources**

GET `/resource-tags`

**Introduced In:** 9.13

Retrieves the tags currently being used for resources in the API.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| num_resources | integer | query | False | Filter by num_resources<br>• Min value: 1 |
| value | string | query | False | Filter by value<br>• maxLength: 200 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Min value: 0<br>• Default value: 15<br>• Max value: 120 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[resource_tag] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "value": "team:accounting"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

resource_tag

A resource tag is a way to group resources in the API together for identification or tracking purposes.

| Name | Type | Description |
|---|---|---|
| num_resources | integer | The number of resources that are currently using this tag. |
| value | string | A key:value formatted string representing the tag's name. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage resource tags

## Manage resource tags

### Overview

You can use this endpoint to list all of the resources in the API that have been tagged with a particular tag value and to create new tags on existing resources. Each resource has a limit of 64 tags that cannot be exceeded. See the section labelled Tagging Resources for Tracking Purposes to find out more information about how to tag a resource.

#### Examples

The following examples show some ways that this endpoint can be used.

**List all resources that are tagged for the test environment.**

```
# The API:
/api/resource-tags/environment:test/resources

# The call:
curl -X GET "https://<mgmt-ip>/api/resource-
tags/environment:test/resources"

# The response:
{
"records": [
  {
    "href": "/api/storage/volumes/558949d1-d4cf-445e-ada5-e340dee6a581",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fstorage%2Fvolumes%2F558949d1-
d4cf-445e-ada5-e340dee6a581"
      }
    }
  },
  {
    "href": "/api/svm/svms/7f97a0b1-fe4f-11e8-b9c5-005056a76061",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fsvm%2Fsvms%2F7f97a0b1-fe4f-11e8-
b9c5-005056a76061"
      }
    }
  },
  {
```

```
    "href": "/api/cluster",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fcluster"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/resource-tags/environment:test/resources"
  }
}
}
```

#### List all volumes that have been tagged for the accounting team

```
# The API:
/api/resource-tags/team:accounting/resources

# The call:
curl -X GET "https://<mgmt-ip>/api/resource-
tags/team:accounting/resources?label=storage_volumes"

# The response:
{
"records": [
  {
    "href": "/api/storage/volumes/558949d1-d4cf-445e-ada5-e340dee6a581",
    "label": "storage_volumes",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fstorage%2Fvolumes%2F558949d1-
d4cf-445e-ada5-e340dee6a581"
      }
    }
  },
  {
    "href": "/api/storage/volumes/64750961-fda7-4327-9f16-00034c3f5ad2",
    "label": "storage_volumes",
    "_links": {
```

```
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fstorage%2Fvolumes%2F64750961-
fda7-4327-9f16-00034c3f5ad2"
      }
    }
  },
  {
    "href": "/api/storage/volumes/bee17b91-f90a-4854-b146-8b102a0a9882",
    "label": "storage_volumes",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/environment:test/resources/%2Fapi%2Fstorage%2Fvolumes%2Fbee17b91-
f90a-4854-b146-8b102a0a9882"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/resource-
tags/environment:test/resources?label=storage_volumes"
  }
}
}
```


### Create a new tag on a volume

```
# The API:
/api/resource-tags/team:accounting/resources

# The call:
curl -X POST "https://<mgmt-ip>/api/resource-
tags/team:accounting/resources" -d '{"href":
"/api/storage/volumes/f288168e-bd3e-11ed-9516-005056acd4e8"}'

# The response
{
"num_records": 1,
"records": [
  {
    "href": "/api/storage/volumes/f288168e-bd3e-11ed-9516-005056acd4e8",
    "_links": {
      "self": {
        "href": "/api/resource-
tags/team:accounting/resources/%2Fapi%2Fstorage%2Fvolumes%2Ff288168e-bd3e-
11ed-9516-005056acd4e8"
      }
    }
  }
]
}
```

---

#### Delete a new tag on a volume

```
# The API:
/resource-tags/{resource_tag.value}/resources/{href}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/resource-
tags/team:accounting/resources/%2Fapi%2Fstorage%2Fvolumes%2Ff288168e-bd3e-
11ed-9516-005056acd4e8"
```

#### Retrieve the resources for a tag

GET `/resource-tags/{resource_tag.value}/resources`

**Introduced In:** 9.13

Retrieves the resources for a specific tag

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| resource_tag.value | string | path | True | • format: key:value |
| value | string | query | False | Filter by value<br>• Introduced in: 9.16 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| href | string | query | False | Filter by href<br>• Introduced in: 9.16 |
| label | string | query | False | Filter by label |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[resource_tag_resource] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "href": "string",
      "label": "volume",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "value": "string"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs
to the cluster as a whole), then this property will be empty and unreturned.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

resource_tag_resource

This object provides a pointer to the tagged resource in the API. Details about the tagged object are available by querying
the address of the href property.

| Name | Type | Description |
| --- | --- | --- |
| href | string | This property provides the address in the API at which the tagged resource is available. Additional queries can be made on this endpoint to fetch the resource's properties. |
| label | string | This is a human-readable classifier representing the type of thing that is pointed to by the href. |
| svm | svm | If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs to the cluster as a whole), then this property will be empty and unreturned. |
| value | string | The text value of the tag formatted as `key:value`. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a new tag for a resource

POST `/resource-tags/{resource_tag.value}/resources`

**Introduced In:** 9.13

Creates a new tag on a specific resource.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| resource_tag.value | string | path | True | • format: key:value |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| href | string | This property provides the address in the API at which the tagged resource is available. Additional queries can be made on this endpoint to fetch the resource's properties. |
| label | string | This is a human-readable classifier representing the type of thing that is pointed to by the href. |
| svm | svm | If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs to the cluster as a whole), then this property will be empty and unreturned. |
| value | string | The text value of the tag formatted as `key:value`. |

### Example request

```
{
  "href": "string",
  "label": "volume",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "value": "string"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262257 | The API specified in the href cannot be found. |
| 262258 | The resource specified in the href does not support tagging. |
| 262259 | The specified href cannot be parsed. |
| 262260 | The provided href is missing a key to a specific resource. |
| 262261 | The specified href points to a resource that does not exist. |
| 262262 | The specified href cannot be an empty string. |
| 262263 | Tag key-value pairs cannot be longer than 200 characters. |

| Error Code | Description |
| --- | --- |
| 263148 | Specified resource cannot be tagged because it reached allowed tags limit. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs
to the cluster as a whole), then this property will be empty and unreturned.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

resource_tag_resource

This object provides a pointer to the tagged resource in the API. Details about the tagged object are available by querying
the address of the href property.

| Name | Type | Description |
|---|---|---|
| href | string | This property provides the address in the API at which the tagged resource is available. Additional queries can be made on this<br>endpoint to fetch the resource's properties. |
| label | string | This is a human-readable classifier representing the type of thing that is pointed to by the href. |

| Name | Type | Description |
|---|---|---|
| svm | svm | If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs to the cluster as a whole), then this property will be empty and unreturned. |
| value | string | The text value of the tag formatted as `key:value`. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a tag for a resource

DELETE `/resource-tags/{resource_tag.value}/resources/{href}`

**Introduced In:** 9.13

Deletes a specific tag on a specific resource.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| resource_tag.value | string | path | True | • format: key:value |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| href | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a resource for a tag

GET `/resource-tags/{resource_tag.value}/resources/{href}`

**Introduced In:** 9.13

Retrieves a specific resource for a specific tag.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| resource_tag.value | string | path | True | • format: key:value |
| href | string | path | True | |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| href | string | This property provides the address in the API at which the tagged resource is available. Additional queries can be made on this endpoint to fetch the resource's properties. |
| label | string | This is a human-readable classifier representing the type of thing that is pointed to by the href. |
| svm | svm | If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs to the cluster as a whole), then this property will be empty and unreturned. |
| value | string | The text value of the tag formatted as `key:value`. |

**Example response**

```
{
  "href": "string",
  "label": "volume",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "value": "string"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

If the tagged resource belongs to an SVM, this property will be set. If the resource does not belong to an SVM (i.e. it belongs
to the cluster as a whole), then this property will be empty and unreturned.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
|---|---|---|
| target | string | The target parameter that caused the error. |

### Retrieve a resource tag

GET `/resource-tags/{value}`

**Introduced In:** 9.13

Retrieves a specific resource tag.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| value | string | path | True | • format: key:value |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| num_resources | integer | The number of resources that are currently using this tag. |
| value | string | A key:value formatted string representing the tag's name. |

#### Example response

```
{
  "value": "team:accounting"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# NAS

## ONTAP REST API NAS endpoints

### Overview

These APIs allow you to complete various tasks, including:

- Creating an NFS server for an SVM
- Managing an NFS configuration of an SVM
- Viewing and updating the NFS configuration of an SVM
- Configuring export policies and rules for an SVM
- Managing export policies and rules for an SVM

### APIs

#### NFS

The NFS APIs enable you to create and configure NFS settings for an SVM. You can delete or update NFS configurations, and you can also disable or enable different NFS features as needed.
The export APIs allow you to create and manage export policies for an SVM that enable an administrator to restrict access to volumes for clients that match specific IP addresses and specific authentication types. Export APIs are also used to create export rules for an export policy. The APIs allow each rule to specify the number of mask bits in the client IP address that must be matched for that rule to apply to a particular client request.

The APIs also allow each export rule to specify the authentication types that are required for both read-only and read-write operations.

#### Kerberos

Kerberos is a protocol designed to provide strong authentication for users and hosts within a client/server environment. The basis of the protocol is a shared, secret-key cryptology system. (Kerberos uses shared-key encryption to ensure the confidentiality of the data. It also uses hashing techniques to ensure the integrity of the data (so that no one can modify the data unless allowed to do so). With the NetApp multiprotocol storage platform, through which clients based on UNIX or Windows can access data using CIFS or NFS, it is crucial to provide the ability to use standard network services for authentication and for identity storage.

To configure an ONTAP system to use Kerberos for NFS, Kerberos must be enabled on a data LIF in the SVM that owns the NFS server. A Kerberos realm needs to be created before enabling Kerberos on a data LIF. (The Kerberos realm is needed so that the cluster knows how to format Kerberos ticket requests.) The Kerberos APIs allow you to define, create, modify, and delete realms for the SVM. The APIs also allow you to enable/disable Kerberos on a data LIF and update the Kerberos interface configuration for a particular data LIF in the SVM.

## Manage Active Directory accounts

### Manage Active Directory accounts

#### Overview

You can use this API to display Active Directory account-related information of all SVMs or create a new Active Directory account.

#### Examples

### Retrieving all Active Directory account details of all SVMs

```
# The API:
GET /api/protocols/active-directory

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/active-directory?fields=*" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "9166c7e8-1fcf-11ed-b037-0050568e9b8b",
      "name": "vs1"
    },
    "organizational_unit": "CN=Computers",
    "name": "SERVER1",
    "fqdn": "EXAMPLE.COM",
    "security": {
      "advertised_kdc_encryptions": [
        "des"
      ]
    }
  }
],
"num_records": 1
}
```

### Create a new Active Directory account

```
# The API:
POST /api/protocols/active-directory

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/active-directory" -H
"accept: application/json"
            -d "{  \"force_account_overwrite\": false,  \"fqdn\":
\"example.com\",  \"name\": \"account1\", \"password\": \"password\",
\"svm\": {    \"_links\": {},    \"name\": \"vs1\",    \"uuid\":
\"6fcfe07a-8a93-11ec-927a-0050568e205f\"  },  \"username\":
\"administrator\", \"security\": { \"advertised_kdc_encryptions\": [
\"des\" ] } }"
```

---

### Retrieve Active Directory accounts for all SVMs

GET `/protocols/active-directory`

**Introduced In:** 9.12

Retrieves Active Directory accounts for all SVMs.

#### Related ONTAP commands

- `vserver active-directory show`
- `vserver active-directory preferred-dc show`
- `vserver active-directory discovered-servers show`
- `vserver cifs security show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| preferred_dcs.server_ip | string | query | False | Filter by preferred_dcs.server_ip |
| preferred_dcs.fqdn | string | query | False | Filter by preferred_dcs.fqdn<br>• maxLength: 254 |
| fqdn | string | query | False | Filter by fqdn<br>• maxLength: 254 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security.advertised_kdc_encryptions | string | query | False | Filter by security.advertised_kdc_encryptions<br>• Introduced in: 9.17 |
| discovered_servers.server.type | string | query | False | Filter by discovered_servers.server.type |
| discovered_servers.server.name | string | query | False | Filter by discovered_servers.server.name |
| discovered_servers.server.ip | string | query | False | Filter by discovered_servers.server.ip |
| discovered_servers.state | string | query | False | Filter by discovered_servers.state |
| discovered_servers.domain | string | query | False | Filter by discovered_servers.domain |
| discovered_servers.node.name | string | query | False | Filter by discovered_servers.node.name |
| discovered_servers.node.uuid | string | query | False | Filter by discovered_servers.node.uuid |
| discovered_servers.preference | string | query | False | Filter by discovered_servers.preference |
| organizational_unit | string | query | False | Filter by organizational_unit |
| name | string | query | False | Filter by name<br>• maxLength: 15 |
| svm.name | string | query | False | Filter by svm.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of Active Directory records. |
| records | array[active_directory] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "discovered_servers": [
        {
          "domain": "server1.com",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "preference": "preferred",
          "server": {
            "ip": "7.4.3.1",
            "name": "scspk0659102001",
            "type": "kerberos"
          },
          "state": "ok"
        }
      ],
      "force_account_overwrite": "",
      "fqdn": "server1.com",
      "name": "account1",
      "organizational_unit": "CN=Test",
      "preferred_dcs": [
        {
          "fqdn": "test.com",
          "server_ip": "4.4.4.4"
        }
      ],
      "security": {
```

```
        "advertised_kdc_encryptions": [
          "string"
        ]
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

active_directory_discovered_server_reference

| Name | Type | Description |
|---|---|---|
| ip | string | The IP address of the server that was discovered. |
| name | string | Server Name |
| type | string | The type of the server that was discovered. |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | The Active Directory domain that the discovered server is a member of. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| preference | string | The preference level of the server that was discovered. |
| server | active_directory_discovered_server_reference | |
| state | string | The status of the connection to the server that was discovered. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

active_directory_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

active_directory

| Name | Type | Description |
| --- | --- | --- |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory (AD) account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account will be created. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an Active Directory account for an SVM**

POST `/protocols/active-directory`

**Introduced In:** 9.12

Creates an Active Directory account for a given SVM.

#### Related ONTAP commands

- `vserver active-directory create`

#### Important notes

- Active Directory account creation is asynchronous. Response contains Task UUID and Link that can be queried to get the status.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory (AD) account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account will be created. |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

**Example request**

```
{
  "discovered_servers": [
    {
      "domain": "server1.com",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "preferred",
      "server": {
        "ip": "7.4.3.1",
        "name": "scspk0659102001",
        "type": "kerberos"
      },
      "state": "ok"
    }
  ],
  "force_account_overwrite": "",
  "fqdn": "server1.com",
  "name": "account1",
  "organizational_unit": "CN=Test",
  "password": "testpwd",
  "preferred_dcs": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ],
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "username": "admin"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory (AD) account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account will be created. |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

**Example response**


```
{
  "discovered_servers": [
    {
      "domain": "server1.com",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "preferred",
      "server": {
        "ip": "7.4.3.1",
        "name": "scspk0659102001",
        "type": "kerberos"
      },
      "state": "ok"
    }
  ],
  "force_account_overwrite": "",
  "fqdn": "server1.com",
  "name": "account1",
  "organizational_unit": "CN=Test",
  "password": "testpwd",
  "preferred_dcs": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ],
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "username": "admin"
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655562 | The NetBIOS name cannot be longer than 15 characters. |
| 655915 | A CIFS server for this SVM already exists. Having both a CIFS server and an Active Directory account for the same SVM is not supported. Use the \"vserver cifs delete\" command to delete the existing CIFS server and try the command again. |
| 656464 | Failed to create the Active Directory machine account. Reason: Invalid Credentials. |
| 656465 | Failed to create the Active Directory machine account. Reason: An account with this name already exists. |
| 656466 | Failed to create the Active Directory machine account. Reason: Unable to connect to any domain controllers. |
| 656467 | Failed to create the Active Directory machine account. Reason: Organizational-Unit not found. |
| 656490 | Unable to create the Active Directory account. The Active Directory account name is already used by another SVM. |
| 656491 | Failed to create the Active Directory machine account. Reason: This operation is supported only on data SVMs. |
| 656494 | The Active Directory was created but was not fully configured. |
| 656495 | Failed to update Active Directory security settings for SVM. |
| 656496 | Failed to update Active Directory security settings for SVM. Reason: KDC has no support for encryption type. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

active_directory_discovered_server_reference

| Name | Type | Description |
|---|---|---|
| ip | string | The IP address of the server that was discovered. |
| name | string | Server Name |
| type | string | The type of the server that was discovered. |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | The Active Directory domain that the discovered server is a member of. |
| node | node | |
| preference | string | The preference level of the server that was discovered. |
| server | active_directory_discovered_server_reference | |
| state | string | The status of the connection to the server that was discovered. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

active_directory_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

active_directory

| Name | Type | Description |
| --- | --- | --- |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory (AD) account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account will be created. |

| Name | Type | Description |
| --- | --- | --- |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage Active Directory accounts for an SVM

### Manage Active Directory accounts for an SVM

#### Overview

You can use this API to display or modify Active Directory account details of the specified SVM.

It can also be used to delete an Active Directory account for the specified SVM.

#### Examples

### Retrieving all Active Directory account details of a specific SVM

```
# The API:
GET /api/protocols/active-directory/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/active-directory/6dd78167-
c907-11eb-b2bf-0050568e7324" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "6dd78167-c907-11eb-b2bf-0050568e7324",
  "name": "vs1"
},
"organizational_unit": "CN=Computers",
"name": "ACCOUNT1",
"fqdn": "EXAMPLE.COM",
"security": {
  "advertised_kdc_encryptions": [
    "des"
  ]
}
}
```

### Update the Active Directory account details of a specific SVM

```
# The API:
PATCH /api/protocols/active-directory/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/active-directory/6dd78167-
c907-11eb-b2bf-0050568e7324" -H "accept: application/json"
            -d "{  \"force_account_overwrite\": true,  \"fqdn\":
\"example.com\",  \"password\": \"password\",  \"username\":
\"administrator\", \"security\": { \"advertised_kdc_encryptions\": [
\"des\" ] } }"
```

---

#### Delete an Active Directory account of a specific SVM

---

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/6dd78167-
c907-11eb-b2bf-0050568e7324" -H "accept: application/json" -d "{
\"password\": \"password\",  \"username\": \"administrator\"}"
```

---

#### Delete an Active Directory account for an SVM

DELETE `/protocols/active-directory/{svm.uuid}`

**Introduced In:** 9.12

Deletes the Active Directory account for a given SVM.

##### Related ONTAP commands

- `vserver active-directory delete`

##### Important notes

- Active Directory account deletion is asynchronous. Response contains Task UUID and Link that can be queried to get the status.

##### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| password | string | Administrator password required for Active Directory account creation. |
| username | string | Administrator username required for Active Directory account creation. |

#### Example request

```
{
  "password": "testpwd",
  "username": "admin"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655414 | Failed to remove Active Directory machine. |
| 655431 | Username cannot be empty. |
| 655432 | Password cannot be empty. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

active_directory_delete

| Name | Type | Description |
|---|---|---|
| password | string | Administrator password required for Active Directory account creation. |
| username | string | Administrator username required for Active Directory account creation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an Active Directory account for an SVM

GET `/protocols/active-directory/{svm.uuid}`

**Introduced In:** 9.12

Retrieves the Active Directory account for a given SVM.

#### Related ONTAP commands

- `vserver active-directory show`
- `vserver active-directory preferred-dc show`

- `vserver active-directory discovered-servers show`
- `vserver active-directory discovered-servers reset-servers`
- `vserver cifs security show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| reset_discovered_servers | boolean | query | False | Force a rediscovery.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| fqdn | string | Fully qualified domain name. |
| name | string | Active Directory (AD) account NetBIOS name. |
| organizational_unit | string | Organizational unit under which the Active Directory account will be created. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "discovered_servers": [
    {
      "domain": "server1.com",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "preferred",
      "server": {
        "ip": "7.4.3.1",
        "name": "scspk0659102001",
        "type": "kerberos"
      },
      "state": "ok"
    }
  ],
  "force_account_overwrite": "",
  "fqdn": "server1.com",
  "name": "account1",
  "organizational_unit": "CN=Test",
  "preferred_dcs": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ],
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ]
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

active_directory_discovered_server_reference

| Name | Type | Description |
|---|---|---|
| ip | string | The IP address of the server that was discovered. |
| name | string | Server Name |
| type | string | The type of the server that was discovered. |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | The Active Directory domain that the discovered server is a member of. |
| node | node | |
| preference | string | The preference level of the server that was discovered. |
| server | active_directory_discovered_server_reference | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The status of the connection to the server that was discovered. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

active_directory_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Modify an Active Directory account for an SVM

PATCH `/protocols/active-directory/{svm.uuid}`

**Introduced In:** 9.12

Modifies the Active Directory account for a given SVM.

#### Related ONTAP commands

- `vserver active-directory modify`
- `vserver cifs security modify`

#### Important notes

- Patching Active Directory account is asynchronous. Response contains Task UUID and Link that can be queried to get the status.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |

| Name | Type | Description |
| --- | --- | --- |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

#### Example request

```
{
  "discovered_servers": [
    {
      "domain": "server1.com",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "preferred",
      "server": {
        "ip": "7.4.3.1",
        "name": "scspk0659102001",
        "type": "kerberos"
      },
      "state": "ok"
    }
  ],
  "force_account_overwrite": "",
  "fqdn": "server1.com",
  "password": "testpwd",
  "preferred_dcs": [
    {}
  ],
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ]
  },
  "username": "admin"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |

| Name | Type | Description |
| --- | --- | --- |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

#### Example response

```
{
  "discovered_servers": [
    {
      "domain": "server1.com",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "preferred",
      "server": {
        "ip": "7.4.3.1",
        "name": "scspk0659102001",
        "type": "kerberos"
      },
      "state": "ok"
    }
  ],
  "force_account_overwrite": "",
  "fqdn": "server1.com",
  "password": "testpwd",
  "preferred_dcs": [
    {}
  ],
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ]
  },
  "username": "admin"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 656464 | Failed to modify the Active Directory machine account. Reason: Invalid Credentials. |

| Error Code | Description |
| --- | --- |
| 656465 | Failed to create the Active Directory machine account. Reason: An account with this name already exists. |
| 656466 | Failed to create the Active Directory machine account. Reason: Unable to connect to any domain controllers. |
| 656467 | Failed to modify the Active Directory machine account. Reason: Organizational-Unit not found. |
| 656478 | Failed to create the Active Directory machine account. Reason: KDC has no support for encryption type. |
| 656490 | Unable to modify the Active Directory account. The Active Directory account name is already used by another SVM. |
| 656495 | Failed to update Active Directory security settings for SVM. |
| 656496 | Failed to update Active Directory security settings for SVM. Reason: KDC has no support for encryption type. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

active_directory_discovered_server_reference

| Name | Type | Description |
|---|---|---|
| ip | string | The IP address of the server that was discovered. |
| name | string | Server Name |
| type | string | The type of the server that was discovered. |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | The Active Directory domain that the discovered server is a member of. |
| node | node | |
| preference | string | The preference level of the server that was discovered. |
| server | active_directory_discovered_server_reference | |
| state | string | The status of the connection to the server that was discovered. |

preferred_dcs

active_directory_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

active_directory

| Name | Type | Description |
| --- | --- | --- |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| force_account_overwrite | boolean | If set to true and a machine account exists with the same name as specified in "name" in Active Directory, it will be overwritten and reused. |
| fqdn | string | Fully qualified domain name. |
| password | string | Administrator password required for Active Directory account creation, modification and deletion. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred domain controller (DC) records. |
| security | active_directory_security | |
| username | string | Administrator username required for Active Directory account creation, modification and deletion. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage Active Directory domain controllers

### Manage Active Directory domain controllers

#### Overview

You can use this API to display the preferred DC configuration of an SVM.

#### Retrieving all the preferred DC configurations of an SVM

---

The preferred DC GET endpoint retrieves all the configurations for a specific SVM.

#### Examples

### Retrieving all fields for all the preferred DC configurations of an SVM

---

```
# The API:
GET /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fields=*" -H
"accept: application/json"

# The response:
{
 "records": [
{
  "svm": {
    "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
    "name": "svm1"
  },
  "fqdn": "host1",
  "server_ip": "4.4.4.4"
},
{
  "svm": {
    "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
    "name": "svm1"
  },
  "fqdn": "host2",
  "server_ip": "11.11.11.11"
}
],
 "num_records": 2
}
```

---

**Retrieving the preferred DC configuration of a specific SVM, "fqdn" and "server_ip"**

---

```
# The API:
GET /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers/{fqdn}/{server_ip}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/host1/4.4.4.4" -H
"accept: application/json"

# The response:
{
"svm": {
  "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
  "name": "svm1"
},
"fqdn": "host1",
"server_ip": "4.4.4.4"
}
```

#### Creating a new preferred DC configuration

The preferred DC POST endpoint creates a new configuration. Both bulk and instance POST is supported.

#### Examples

### Creating a single preferred DC configuration

```
# The API:
POST /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-
controllers?skip_config_validation=true&return_records=false" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fqdn\":
\"testing.com\", \"server_ip\": \"1.1.1.1\"}"
```

#### Deleting an existing preferred DC configuration

The preferred DC DELETE endpoint deletes an existing configuration. Both bulk and instance delete is supported.

#### Examples

### Deleting the preferred DC configuration of a specific SVM, "fqdn" and "server_ip"

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers/{fqdn}/{server_ip}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/sample/4.4.4.4"
```

### Deleting the preferred DC configurations of a specific SVM and "fqdn"

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fqdn=sampl"
```

### Deleting all preferred DC configurations of a specific SVM

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fqdn=*"
```

#### Deleting the preferred DC configurations of a specific SVM, "fqdn" and set of "server_ips"

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-
controllers/?fqdn=sample&server_ip=3.3.3.3&#124;4.4.4.4&#124;1.1.1.1&#124;
2.2.2.2"
```

#### Deleting the preferred DC configurations of a specific SVM and set of "server_ips"

```
# The API:
DELETE /api/protocols/active-directory/{svm.uuid}/preferred-domain-
controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/active-directory/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-
controllers/?server_ip=3.3.3.3&#124;4.4.4.4&#124;1.1.1.1&#124;2.2.2.2"
```

### Retrieve an Active Directory preferred DC configuration for an SVM

GET `/protocols/active-directory/{svm.uuid}/preferred-domain-controllers`

**Introduced In:** 9.12

Retrieves the Active Directory preferred DC configuration of an SVM.

#### Related ONTAP commands

- `vserver active-directory preferred-dc show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| server_ip | string | query | False | Filter by server_ip |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fqdn | string | query | False | Filter by fqdn<br>• maxLength: 254 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of Active Directory preferred DC records. |
| records | array[active_directory_preferred_dc] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

active_directory_preferred_dc

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an Active Directory preferred DC configuration for an SVM

POST `/protocols/active-directory/{svm.uuid}/preferred-domain-controllers`

**Introduced In:** 9.12

Creates an Active Directory preferred DC configuration for an SVM.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the preferred DC.
- `fqdn` - Fully Qualified Domain Name.
- `server_ip` - IPv4/IPv6 address of the preferred DC.

#### The following parameters are optional:

- skip_config_validation

#### Related ONTAP commands

- `vserver active-directory preferred-dc add`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| skip_config_validation | boolean | query | False | Skip the validation of the specified preferred DC configuration.<br>• Default value: |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |

| Name | Type | Description |
|---|---|---|
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

### Example request

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

### Example response

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4"
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655366 | Invalid domain controller. |
| 655506 | Failed to add preferred-dc. |
| 655918 | The fully qualified domain name cannot be longer than 254 bytes. |
| 656407 | Failed to validate preferred-dc. |
| 656408 | RPC failure occurred during preferred-dc config validation. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

active_directory_preferred_dc

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an Active Directory preferred DC configuration for an SVM and domain

DELETE `/protocols/active-directory/{svm.uuid}/preferred-domain-controllers/{fqdn}/{server_ip}`

**Introduced In:** 9.12

Deletes the Active Directory preferred DC configuration of the specified SVM and domain.

#### Related ONTAP commands

- `vserver active-directory preferred-dc delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fqdn | string | path | True | Fully Qualified Domain Name. |
| server_ip | string | path | True | DC's IP address. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655507 | Failed to remove preferred-dc. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an Active Directory preferred DC configuration for an SVM and domain

GET `/protocols/active-directory/{svm.uuid}/preferred-domain-controllers/{fqdn}/{server_ip}`

**Introduced In:** 9.12

Retrieves the Active Directory preferred DC configuration of an SVM.

#### Related ONTAP commands

- `vserver active-directory preferred-dc show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fqdn | string | path | True | Fully Qualified Domain Name. |
| server_ip | string | path | True | DC's IP address. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred DC. The address can be either an IPv4 or an IPv6 address. |

#### Example response

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NAS audit configurations

### Manage NAS audit configurations

#### Overview

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on storage virtual machines (SVMs). This helps you track potential security problems and provides evidence of any security breaches.

#### Examples

---

### Creating an audit entry with log rotation size and log retention count

To create an audit entry with log rotation size and log retention count, use the following API. Note the *return_records=true* query parameter is used to obtain the newly created entry in the response.

---

```
# The API:
POST /api/protocols/audit/

# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/protocols/audit" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
true, \"events\": { \"async_delete\": false, \"audit_policy_change\":
true, \"authorization_policy\": false, \"cap_staging\": false,
\"cifs_logon_logoff\": true, \"file_operations\": true, \"file_share\":
false, \"security_group\": false, \"user_account\": false },
\"guarantee\": true, \"log\": { \"format\": \"evtx\", \"retention\": {
\"count\": 10 }, \"rotation\": { \"size\": 2048000 }}, \"log_path\":
\"/\", \"svm\": { \"name\": \"vs1\", \"uuid\": \"ec650e97-156e-11e9-abcb-
005056bbd0bf\" }}"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "enabled": true,
    "events": {
      "authorization_policy": false,
      "cap_staging": false,
      "cifs_logon_logoff": true,
      "file_operations": true,
      "file_share": false,
      "security_group": false,
      "user_account": false,
      "audit_policy_change": true,
      "async_delete": false
    },
    "log": {
      "format": "evtx",
      "rotation": {
        "size": 2048000
      },
      "retention": {
        "count": 10,
        "duration": "0s"
      }
    },
    "log_path": "/",
    "guarantee": true
  }
],
"num_records": 1
```

```
}
```

### Creating an audit entry with log rotation schedule and log retention duration

To create an audit entry with log rotation schedule and log retention duration, use the following API. Note that the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/audit/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/audit" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"events\": { \"async_delete\": false, \"audit_policy_change\":
true, \"authorization_policy\": false, \"cap_staging\": false,
\"cifs_logon_logoff\": true, \"file_operations\": true, \"file_share\":
false, \"security_group\": false, \"user_account\": false },
\"guarantee\": true, \"log\": { \"format\": \"xml\", \"retention\": {
\"duration\": \"P4DT12H30M5S\" }, \"rotation\": { \"schedule\": {
\"days\": [1, 5, 10, 15], \"hours\": [0, 1, 6, 12, 18, 23], \"minutes\":
[10, 15, 30, 45, 59], \"months\": [0], \"weekdays\": [0, 2, 5] } } },
\"log_path\": \"/\", \"svm\": { \"name\": \"vs3\", \"uuid\": \"a8d64674-
13fc-11e9-87b1-005056a7ae7e\" }}"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "enabled": true,
    "events": {
      "authorization_policy": false,
      "cap_staging": false,
      "cifs_logon_logoff": true,
      "file_operations": true,
      "file_share": false,
      "security_group": false,
      "user_account": false,
      "audit_policy_change": true,
      "async_delete": false
```

```
    },
    "log": {
      "format": "xml",
      "rotation": {
        "schedule": {
          "minutes": [
            10,
            15,
            30,
            45,
            59
          ],
          "hours": [
            0,
            1,
            6,
            12,
            18,
            23
          ],
          "weekdays": [
            0,
            2,
            5
          ],
          "days": [
            1,
            5,
            10,
            15
          ],
          "months": [
            0
          ]
        }
      },
      "retention": {
        "count": 0,
        "duration": "P4DT12H30M5S"
      }
    },
    "log_path": "/",
    "guarantee": true
  }
],
"num_records": 1
```

```
}
```

### Retrieving an audit configuration for all SVMs in the cluster

```
# The API:
GET /api/protocols/audit/

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/audit?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "enabled": true,
    "events": {
      "authorization_policy": false,
      "cap_staging": false,
      "cifs_logon_logoff": true,
      "file_operations": true,
      "file_share": false,
      "security_group": false,
      "user_account": false,
      "audit_policy_change": true,
      "async_delete": false
    },
    "log": {
      "format": "evtx",
      "rotation": {
        "size": 2048000
      },
      "retention": {
        "count": 10,
        "duration": "0s"
      }
    },
```

```
    "log_path": "/",
    "guarantee": true,
    "charge_qos": false
  },
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "enabled": true,
    "events": {
      "authorization_policy": false,
      "cap_staging": false,
      "cifs_logon_logoff": true,
      "file_operations": true,
      "file_share": false,
      "security_group": false,
      "user_account": false,
      "audit_policy_change": true,
      "async_delete": false
    },
    "log": {
      "format": "xml",
      "rotation": {
        "schedule": {
          "minutes": [
            10,
            15,
            30,
            45,
            59
          ],
          "hours": [
            0,
            1,
            6,
            12,
            18,
            23
          ],
          "weekdays": [
            0,
            2,
            5
          ],
          "days": [
```

```
            1,
            5,
            10,
            15
          ],
          "months": [
            0
          ]
        }
      },
      "retention": {
        "count": 0,
        "duration": "P4DT12H30M5S"
      }
    },
    "log_path": "/",
    "guarantee": true,
    "charge_qos": false
  }
],
"num_records": 2
}
```


---

**Retrieving specific entries with event list as cifs-logon-logoff, file-ops = true for an SVM**

The configuration returned is identified by the events in the list of audit configurations for an SVM.

---

```
# The API:
GET /api/protocols/audit/

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/audit?events.file_operations=true&events.cifs_logon_logo
ff=true&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "events": {
      "cifs_logon_logoff": true,
      "file_operations": true
    }
  },
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "events": {
      "cifs_logon_logoff": true,
      "file_operations": true
    }
  }
],
"num_records": 2
}
```

#### Retrieving a specific audit configuration for an SVM

The configuration returned is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/audit/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
  "name": "vs1"
},
"enabled": true,
"events": {
  "authorization_policy": false,
  "cap_staging": false,
  "cifs_logon_logoff": true,
  "file_operations": true,
  "file_share" : false,
  "security_group": false,
  "user_account": false,
  "audit_policy_change": true,
  "async_delete": false
},
"log": {
  "format": "evtx",
  "rotation": {
    "size": 2048000
  },
  "retention": {
    "count": 10,
    "duration": "0s"
  }
},
"log_path": "/",
"guarantee": true,
"charge_qos": false
}
```

#### Updating a specific audit configuration of an SVM

The configuration is identified by the UUID of its SVM and the provided information is updated.

```
# The API:
PATCH /api/protocols/audit/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"enabled\": false}"

# The response:
{
"job": {
  "uuid": "9938d743-d566-11eb-ad60-0050568e5902",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9938d743-d566-11eb-ad60-0050568e5902"
    }
  }
}
}
```

#### Deleting a specific audit configuration for an SVM

The entry to be deleted is identified by the UUID of its SVM.

```
# The API:
DELETE /api/protocols/audit/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf?force=true" -H "accept: application/json"

# The response:
{
"job": {
  "uuid": "9938d743-d566-11eb-ad60-0050568e5902",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9938d743-d566-11eb-ad60-0050568e5902"
    }
  }
}
}
```

### Retrieve audit configurations

GET `/protocols/audit`

**Introduced In:** 9.6

Retrieves audit configurations.

#### Related ONTAP commands

- `vserver audit show`

#### Learn more

- DOC /protocols/audit

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| log.format | string | query | False | Filter by log.format |
| log.rotation.size | integer | query | False | Filter by log.rotation.size |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| log.rotation.schedule.weekdays | integer | query | False | Filter by log.rotation.schedule.weekdays<br>• Max value: 6<br>• Min value: 0 |
| log.rotation.schedule.minutes | integer | query | False | Filter by log.rotation.schedule.minutes<br>• Max value: 59<br>• Min value: 0 |
| log.rotation.schedule.days | integer | query | False | Filter by log.rotation.schedule.days<br>• Max value: 31<br>• Min value: 1 |
| log.rotation.schedule.months | integer | query | False | Filter by log.rotation.schedule.months<br>• Max value: 12<br>• Min value: 1 |
| log.rotation.schedule.hours | integer | query | False | Filter by log.rotation.schedule.hours<br>• Max value: 23<br>• Min value: 0 |
| log.retention.count | integer | query | False | Filter by log.retention.count |
| log.retention.duration | string | query | False | Filter by log.retention.duration |
| guarantee | boolean | query | False | Filter by guarantee<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| charge_qos | boolean | query | False | Filter by charge_qos<br>• Introduced in: 9.16 |
| enabled | boolean | query | False | Filter by enabled |
| log_path | string | query | False | Filter by log_path |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| events.security_group | boolean | query | False | Filter by events.security_group |
| events.authorization_policy | boolean | query | False | Filter by events.authorization_policy |
| events.async_delete | boolean | query | False | Filter by events.async_delete<br>• Introduced in: 9.16 |
| events.audit_policy_change | boolean | query | False | Filter by events.audit_policy_change<br>• Introduced in: 9.16 |
| events.cap_staging | boolean | query | False | Filter by events.cap_staging |
| events.user_account | boolean | query | False | Filter by events.user_account |
| events.cifs_logon_logoff | boolean | query | False | Filter by events.cifs_logon_logoff |
| events.file_share | boolean | query | False | Filter by events.file_share |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| events.file_operations | boolean | query | False | Filter by events.file_operations |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[audit] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "charge_qos": "",
      "guarantee": "",
      "log": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "format": "string",
        "retention": {
          "duration": "P4DT12H30M5S"
        },
        "rotation": {
          "schedule": {
            "days": [
              "integer"
            ],
            "hours": [
              "integer"
            ],
            "minutes": [
              "integer"
            ],
            "months": [
              "integer"
            ],
            "weekdays": [
              "integer"
            ]
          }
        }
      },
```

```
      "log_path": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

events

| Name | Type | Description |
|---|---|---|
| async_delete | boolean | Volume file async delete events |
| audit_policy_change | boolean | Audit policy change events |
| authorization_policy | boolean | Authorization policy change events |
| cap_staging | boolean | Central access policy staging events |
| cifs_logon_logoff | boolean | CIFS logon and logoff events |
| file_operations | boolean | File operation events |
| file_share | boolean | File share category events |
| security_group | boolean | Local security group management events |
| user_account | boolean | Local user account management events |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

retention

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration. |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count. |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
| --- | --- | --- |
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

log

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| format | string | The format in which the logs are generated by consolidation process.<br>Possible values are:<br>• xml - Data ONTAP-specific XML log format<br>• evtx - Microsoft Windows EVTX log format<br>◦ Default value: 1<br>◦ enum: ["xml", "evtx"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

| Name | Type | Description |
|---|---|---|
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an audit configuration

POST `/protocols/audit`

**Introduced In:** 9.6

Creates an audit configuration.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM to which audit configuration is to be created.
- `log_path` - Path in the owning SVM namespace that is used to store audit logs.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `events.authorization_policy` - *false*
- `events.cap_staging` - *false*
- `events.file_share` - *false*
- `events.security_group` - *false*
- `events.user_account` - *false*
- `events.cifs_logon_logoff` - *true*
- `events.audit_policy_change` - *true*
- `events.async_delete` - *false*
- `events.file_operations` - *true*
- `log.format` - *evtx*
- `log.retention.count` - *0*
- `log.retention.duration` - *PT0S*
- `log.rotation.size` - *100MB*

- `log.rotation.now` - *false*
- `guarantee` - *true*
- `charge_qos` - *false*

#### Related ONTAP commands

- `vserver audit create`
- `vserver audit enable`

#### Learn more

- DOC /protocols/audit

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "charge_qos": "",
  "guarantee": "",
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262186 | Field "log.retention.duration" cannot be used with field "log.retention.count" |
| 262196 | Log_rotation_now is not an allowed operation |
| 2621462 | The specified SVM does not exist |
| 9699330 | An audit configuration already exists |
| 9699337 | Audit system internal update is in progress, audit configuration create failed |
| 9699340 | SVM UUID lookup failed |
| 9699358 | Audit configuration is absent for enabling |
| 9699359 | Audit configuration is already enabled |
| 9699360 | Final consolidation is in progress, audit enable failed |
| 9699365 | Enabling of audit configuration failed |
| 9699370 | Auditing was successfully configured, however audit configuration could not be enabled |

| Error Code | Description |
| --- | --- |
| 9699384 | The specified log_path does not exist |
| 9699385 | The log_path must be a directory |
| 9699386 | The log_path must be a canonical path in the SVMs namespace |
| 9699387 | The log_path cannot be empty |
| 9699388 | Rotate size must be greater than or equal to 1024 KB |
| 9699389 | The log_path must not contain a symbolic link |
| 9699398 | The log_path exceeds a maximum supported length of characters |
| 9699399 | The log_path contains an unsupported read-only (DP/LS) volume |
| 9699400 | The specified log_path is not a valid destination for SVM |
| 9699402 | The log_path contains an unsupported snaplock volume |
| 9699403 | The log_path cannot be accessed for validation |
| 9699406 | The log_path validation failed |
| 9699409 | Failed to enable multiproto.audit.evtxlog.support support capability |
| 9699428 | All nodes need to run ONTAP 8.3.0 release to audit CIFS logon-logoff events |
| 9699429 | Failed to enable multiproto.audit.cifslogonlogoff.support support capability |
| 9699431 | All nodes need to run ONTAP 8.3.0 release to audit CAP staging events |
| 9699432 | Failed to enable multiproto.audit.capstaging.support support capability |

**Definitions**

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| async_delete | boolean | Volume file async delete events |
| audit_policy_change | boolean | Audit policy change events |
| authorization_policy | boolean | Authorization policy change events |
| cap_staging | boolean | Central access policy staging events |
| cifs_logon_logoff | boolean | CIFS logon and logoff events |
| file_operations | boolean | File operation events |
| file_share | boolean | File share category events |
| security_group | boolean | Local security group management events |
| user_account | boolean | Local user account management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration. |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count. |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
| --- | --- | --- |
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |

| Name | Type | Description |
| --- | --- | --- |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

log

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format in which the logs are generated by consolidation process.<br>Possible values are:<br>• xml - Data ONTAP-specific XML log format<br>• evtx - Microsoft Windows EVTX log format<br>◦ Default value: 1<br>◦ enum: ["xml", "evtx"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

| Name | Type | Description |
| --- | --- | --- |
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an audit configuration

DELETE `/protocols/audit/{svm.uuid}`

**Introduced In:** 9.6

Deletes an audit configuration.

#### Related ONTAP commands

- `vserver audit disable`
- `vserver audit delete`

#### Learn more

- DOC /protocols/audit

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Indicates to force delete audit configuration.<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9699349 | Auditing should be disabled before deleting the audit configuration |
| 9699350 | Audit configuration cannot be deleted, final consolidation is in progress |
| 9699410 | Failed to disable multiproto.audit.evtxlog.support support capability |
| 9699430 | Failed to disable multiproto.audit.cifslogonlogoff.support support capability |
| 9699433 | Failed to disable multiproto.audit.capstaging.support support capability |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve the audit configuration for an SVM**

GET `/protocols/audit/{svm.uuid}`

**Introduced In:** 9.6

Retrieves an audit configuration for an SVM.

#### Related ONTAP commands

- `vserver audit show`

#### Learn more

- DOC /protocols/audit

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response

```
{
  "charge_qos": "",
  "guarantee": "",
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| async_delete | boolean | Volume file async delete events |
| audit_policy_change | boolean | Audit policy change events |
| authorization_policy | boolean | Authorization policy change events |
| cap_staging | boolean | Central access policy staging events |
| cifs_logon_logoff | boolean | CIFS logon and logoff events |
| file_operations | boolean | File operation events |
| file_share | boolean | File share category events |
| security_group | boolean | Local security group management events |
| user_account | boolean | Local user account management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration. |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count. |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
| --- | --- | --- |
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |

| Name | Type | Description |
| --- | --- | --- |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

log

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| format | string | The format in which the logs are generated by consolidation process.<br>Possible values are:<br>• xml - Data ONTAP-specific XML log format<br>• evtx - Microsoft Windows EVTX log format<br>◦ Default value: 1<br>◦ enum: ["xml", "evtx"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the audit configuration for an SVM

PATCH `/protocols/audit/{svm.uuid}`

**Introduced In:** 9.6

Updates an audit configuration for an SVM.

#### Related ONTAP commands

- `vserver audit modify`

#### Learn more

- DOC /protocols/audit

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |

| Name | Type | Description |
| --- | --- | --- |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |

**Example request**

```
{
  "charge_qos": "",
  "guarantee": "",
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262186 | Field "log.retention.duration" cannot be used with field "log.retention.count" |
| 9699340 | SVM UUID lookup failed |
| 9699343 | Audit configuration is absent for modification |
| 9699358 | Audit configuration is absent for enabling |
| 9699359 | Audit configuration is already enabled |
| 9699360 | Final consolidation is in progress, audit enable failed |
| 9699365 | Enabling of audit configuration failed |
| 9699373 | Audit configuration is absent for disabling |
| 9699374 | Audit configuration is already disabled |
| 9699375 | Disabling of audit configuration failed |
| 9699384 | The specified log_path does not exist |

| Error Code | Description |
| --- | --- |
| 9699385 | The log_path must be a directory |
| 9699386 | The log_path must be a canonical path in the SVMs namespace |
| 9699387 | The log_path cannot be empty |
| 9699388 | Rotate size must be greater than or equal to 1024 KB |
| 9699389 | The log_path must not contain a symbolic link |
| 9699398 | The log_path exceeds a maximum supported length of characters |
| 9699399 | The log_path contains an unsupported read-only (DP/LS) volume |
| 9699400 | The specified log_path is not a valid destination for SVM |
| 9699402 | The log_path contains an unsupported snaplock volume |
| 9699403 | The log_path cannot be accessed for validation |
| 9699406 | The log_path validation failed |
| 9699407 | Additional fields are provided |
| 9699409 | Failed to enable multiproto.audit.evtxlog.support support capability |
| 9699410 | Failed to disable multiproto.audit.evtxlog.support support capability |
| 9699418 | Audit configuration is absent for rotate |
| 9699419 | Failed to rotate audit log |
| 9699420 | Cannot rotate audit log, auditing is not enabled for this SVM |
| 9699428 | All nodes need to run ONTAP 8.3.0 release to audit CIFS logon-logoff events |
| 9699429 | Failed to enable multiproto.audit.cifslogonlogoff.support support capability |
| 9699430 | Failed to disable multiproto.audit.cifslogonlogoff.support support capability |
| 9699431 | All nodes need to run ONTAP 8.3.0 release to audit CAP staging events |
| 9699432 | Failed to enable multiproto.audit.capstaging.support support capability |
| 9699433 | Failed to disable multiproto.audit.capstaging.support support capability |

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| async_delete | boolean | Volume file async delete events |
| audit_policy_change | boolean | Audit policy change events |
| authorization_policy | boolean | Authorization policy change events |
| cap_staging | boolean | Central access policy staging events |
| cifs_logon_logoff | boolean | CIFS logon and logoff events |
| file_operations | boolean | File operation events |
| file_share | boolean | File share category events |
| security_group | boolean | Local security group management events |
| user_account | boolean | Local user account management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration. |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count. |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
| --- | --- | --- |
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |

| Name | Type | Description |
| --- | --- | --- |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

log

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format in which the logs are generated by consolidation process.<br>Possible values are:<br>• xml - Data ONTAP-specific XML log format<br>• evtx - Microsoft Windows EVTX log format<br>◦ Default value: 1<br>◦ enum: ["xml", "evtx"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

| Name | Type | Description |
| --- | --- | --- |
| charge_qos | boolean | Indicates if audit logs generation should incur an extra charge. |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| guarantee | boolean | Indicates whether there is a strict Guarantee of Auditing |
| log | log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 audit configurations

### Manage S3 audit configurations

#### Overview

S3 events auditing is a security measure that enables you to track and log certain S3 events on storage virtual machines (SVMs). You can track potential security problems and provides evidence of any security breaches.

#### Examples

---

### Creating an S3 audit entry with log rotation size and log retention count

To create an S3 audit entry with log rotation size and log retention count, use the following API. Note the *return_records=true* query parameter is used to obtain the newly created entry in the response.

---

```
# The API:
POST /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf/object-store?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
true, \"events\": { \"data\": false, \"management\": false}, \"log\": {
\"format\": \"json\", \"retention\": { \"count\": 10 }, \"rotation\": {
\"size\": 2048000 }}, \"log_path\": \"/\"}"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "enabled": true,
    "events": {
      "data": false,
      "management": false
    },
    "log": {
      "format": "json",
      "rotation": {
        "size": 2048000
      },
      "retention": {
        "count": 10,
        "duration": "0s"
      }
    },
    "log_path": "/"
  }
],
"num_records": 1
}
```

#### Creating an S3 audit entry with log rotation schedule and log retention duration

To create an S3 audit entry with log rotation schedule and log retention duration, use the following API. Note that the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/audit/a8d64674-13fc-11e9-
87b1-005056a7ae7e/object-store?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"events\": { \"data\": true, \"management\": true }, \"log\": {
\"format\": \"json\", \"retention\": { \"duration\": \"P4DT12H30M5S\" },
\"rotation\": { \"schedule\": { \"days\": [1, 5, 10, 15], \"hours\": [0,
1, 6, 12, 18, 23], \"minutes\": [10, 15, 30, 45, 59], \"months\": [0],
\"weekdays\": [0, 2, 5] } } }, \"log_path\": \"/\"}"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "enabled": true,
    "events": {
      "data": true,
      "management": true
    },
    "log": {
      "format": "json",
      "rotation": {
        "schedule": {
          "minutes": [
            10,
            15,
            30,
            45,
            59
          ],
          "hours": [
            0,
            1,
            6,
            12,
            18,
            23
          ],
```

```
          "weekdays": [
            0,
            2,
            5
          ],
          "days": [
            1,
            5,
            10,
            15
          ],
          "months": [
            0
          ]
        }
      },
      "retention": {
        "count": 0,
        "duration": "P4DT12H30M5S"
      }
    },
    "log_path": "/"
  }
],
"num_records": 1
}
```

#### Retrieving an S3 audit configuration for all SVMs in the cluster

```
# The API:
GET /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/audit/*/object-
store?fields=*&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
```

```
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "enabled": true,
    "events": {
      "data": false,
      "management": false
    },
    "log": {
      "format": "json",
      "rotation": {
        "size": 2048000
      },
      "retention": {
        "count": 10,
        "duration": "0s"
      }
    },
    "log_path": "/"
  },
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "enabled": true,
    "events": {
      "data": true,
      "management": true
    },
    "log": {
      "format": "json",
      "rotation": {
        "schedule": {
          "minutes": [
            10,
            15,
            30,
            45,
            59
          ],
          "hours": [
            0,
            1,
            6,
            12,
```

```
            18,
            23
          ],
          "weekdays": [
            0,
            2,
            5
          ],
          "days": [
            1,
            5,
            10,
            15
          ],
          "months": [
            0
          ]
        }
      },
      "retention": {
        "count": 0,
        "duration": "P4DT12H30M5S"
      }
    },
    "log_path": "/"
  }
],
"num_records": 2
}
```


#### Retrieving specific entries with event list as data and management event for an SVM

The configuration returned is identified by the events in the list of S3 audit configurations of an SVM.

```
# The API:
GET /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/audit/*/object-
store?events.data=true&events.management=true&return_records=true&return_t
imeout=15" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
      "name": "vs1"
    },
    "events": {
      "data": true,
      "management":true
    }
  },
  {
    "svm": {
      "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
      "name": "vs3"
    },
    "events": {
      "data": true,
      "management": true
    }
  }
],
"num_records": 2
}
```

#### Retrieving a specific S3 audit configuration of an SVM

The configuration returned is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf/object-store/" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
  "name": "vs1"
},
"enabled": true,
"events": {
  "data": false,
  "management": false
},
"log": {
  "format": "json",
  "rotation": {
    "size": 2048000
  },
  "retention": {
    "count": 10,
    "duration": "0s"
  }
},
"log_path": "/"
}
```

#### Updating a specific S3 audit configuration of an SVM

The configuration is identified by the UUID of its SVM and the provided information is updated.

```
# The API:
PATCH /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf/object-store/" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"enabled\": false}"
```

---

### Deleting a specific S3 audit configuration of an SVM

The entry to be deleted is identified by the UUID of its SVM.

---

```
# The API:
DELETE /api/protocols/audit/{svm.uuid}/object-store/

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf/object-store" -H "accept: application/json"
```

### Delete an S3 audit configuration

DELETE `/protocols/audit/{svm.uuid}/object-store`

**Introduced In:** 9.10

Deletes an S3 audit configuration.

#### Related ONTAP commands

- `vserver object-store-server audit disable`
- `vserver object-store-server audit delete`

#### Learn more

- DOC /protocols/audit/{svm.uuid}/object-store

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Indicates whether a force deletion of the audit configuration is enabled. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 140902420 | Failed to delete audit configuration for the SVM. |
| 140902421 | Failed to delete audit configuration for the SVM because audit is enabled for the SVM. |
| 140902422 | Failed to delete audit configuration for the SVM because final consolidation is in progress. Wait a few minutes, and try the operation again. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve an S3 audit configurations**

GET `/protocols/audit/{svm.uuid}/object-store`

**Introduced In:** 9.10

Retrieves S3 audit configurations.

#### Related ONTAP commands

- `vserver object-store-server audit show`

#### Learn more

- DOC /protocols/audit/{svm.uuid}/object-store

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| log_path | string | query | False | Filter by log_path |
| svm.name | string | query | False | Filter by svm.name |
| events.data | boolean | query | False | Filter by events.data |
| events.management | boolean | query | False | Filter by events.management |
| log.format | string | query | False | Filter by log.format |
| log.retention.duration | string | query | False | Filter by log.retention.duration |
| log.retention.count | integer | query | False | Filter by log.retention.count |
| log.rotation.size | integer | query | False | Filter by log.rotation.size |
| log.rotation.schedule.weekdays | integer | query | False | Filter by log.rotation.schedule.weekdays<br>• Max value: 6<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| log.rotation.schedule.minutes | integer | query | False | Filter by log.rotation.schedule.minutes<br>• Max value: 59<br>• Min value: 0 |
| log.rotation.schedule.days | integer | query | False | Filter by log.rotation.schedule.days<br>• Max value: 31<br>• Min value: 1 |
| log.rotation.schedule.months | integer | query | False | Filter by log.rotation.schedule.months<br>• Max value: 12<br>• Min value: 1 |
| log.rotation.schedule.hours | integer | query | False | Filter by log.rotation.schedule.hours<br>• Max value: 23<br>• Min value: 0 |
| enabled | boolean | query | False | Filter by enabled |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| log | s3_log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| data | boolean | Data events |
| management | boolean | Management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with "duration". |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with "count". |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |

| Name | Type | Description |
| --- | --- | --- |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

s3_log

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| format | string | Format in which the logs are generated by the consolidation process.<br>Possible values are:<br>• json - ONTAP-specific Json log format.<br>◦ Default value: 1<br>◦ enum: ["json"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an S3 audit configuration for an SVM

PATCH `/protocols/audit/{svm.uuid}/object-store`

**Introduced In:** 9.10

Updates an S3 audit configuration for an SVM.

#### Important notes

- `events` - Not specifying either data or management is equivalent to setting it to false.

#### Related ONTAP commands

- `vserver object-store-server audit modify`

#### Learn more

- DOC /protocols/audit/{svm.uuid}/object-store

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| log | s3_log | |

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |

**Example request**

```
{
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 140902401 | Failed to create an audit configuration for the SVM. |
| 140902402 | Audit configuration is already present. |
| 140902402 | Audit configuration is already enabled. |
| 140902403 | Failed to create staging volume. |
| 140902415 | Failed to modify an audit configuration because no audit configuration exists for SVM. |
| 140902416 | Failed to modify audit configuration for SVM. |
| 140902422 | Final consolidation is in progress, audit delete failed. |
| 140902423 | Failed to delete the audit configuration for the SVM. |
| 140902425 | Audit configuration is not available for disabling. |
| 140902430 | Audit configuration is not available for enabling. |
| 140902431 | Audit enable failed, audit configuration already enabled for the SVM. |
| 140902432 | Final consolidation is in progress, audit enable failed. |
| 140902445 | Audit disable failed, audit configuration does not exist for the SVM. |
| 140902446 | Audit configuration is already disabled. |
| 140902446 | Audit disable failed, audit configuration does not exist for the SVM. |

| Error Code | Description |
| --- | --- |
| 140902456 | The specified log_path does not exist. |
| 140902457 | The log_path must be a directory. |
| 140902458 | The log_path must be a canonical path in the SVM's namespace. |
| 140902459 | The log_path cannot be empty. |
| 140902460 | Rotate size must be greater than or equal to 1024 KB. |
| 140902461 | The destination path must not contain a symbolic link. |
| 140902470 | The log_path exceeds a maximum supported length of characters. |
| 140902471 | The log_path contains an unsupported read-only (DP/LS) volume. |
| 140902472 | The log_path is not a valid destination for the SVM. |
| 140902474 | The log_path contains an unsupported Snaplock volume. |
| 140902478 | The log_path validation failed. |
| 140902478 | The log_path cannot be accessed for validation. |
| 140902479 | No other fields can be specified when "enabled" is specified. |
| 140902490 | Audit configuration is absent for rotate. |
| 140902491 | Failed to rotate audit log. |
| 140902492 | Cannot rotate audit log, auditing is not enabled for this SVM. |

Also see the table of common errors in the Response body overview section of this documentation.
ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 262186 | Field "log.retention.duration" cannot be used with field "log.retention.count" |
| 9699340 | SVM UUID lookup failed |
| 9699407 | Additional fields are provided |

#### Definitions

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| data | boolean | Data events |
| management | boolean | Management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with "duration". |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with "count". |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |

| Name | Type | Description |
| --- | --- | --- |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

s3_log

| Name | Type | Description |
| --- | --- | --- |
| format | string | Format in which the logs are generated by the consolidation process.<br>Possible values are:<br>• json - ONTAP-specific Json log format.<br>◦ Default value: 1<br>◦ enum: ["json"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_audit

Auditing for NAS events is a security measure that enables you to track and log certain S3 events on SVMs.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| events | events | |
| log | s3_log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an S3 audit configuration**

POST `/protocols/audit/{svm.uuid}/object-store`

**Introduced In:** 9.10

Creates an S3 audit configuration.

#### Required properties

- `log_path` - Path in the owning SVM namespace that is used to store audit logs.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `events.data` - *true*
- `events.management` - *false*
- `log.format` - *json*
- `log.retention.count` - *0*
- `log.retention.duration` - *PT0S*
- `log.rotation.size` - *100MB*
- `log.rotation.now` - *false*

#### Related ONTAP commands

- `vserver object-store-server audit create`
- `vserver object-store-server audit enable`

#### Learn more

- DOC /protocols/audit/{svm.uuid}/object-store

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |
| events | events | |
| log | s3_log | |

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[s3_audit] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "log": {
        "format": "string",
        "retention": {
          "duration": "P4DT12H30M5S"
        },
        "rotation": {
          "schedule": {
            "days": [
              "integer"
            ],
            "hours": [
              "integer"
            ],
            "minutes": [
              "integer"
            ],
            "months": [
              "integer"
            ],
            "weekdays": [
              "integer"
            ]
          }
        }
      },
      "log_path": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 140902401 | Failed to create an audit configuration for the SVM. |
| 140902402 | Audit configuration is already present. |
| 140902402 | Audit configuration is already enabled. |
| 140902403 | Failed to create staging volume. |
| 140902415 | Failed to modify an audit configuration because no audit configuration exists for the SVM. |
| 140902416 | Failed to modify audit configuration for SVM. |
| 140902422 | Final consolidation is in progress, audit delete failed. |
| 140902423 | Failed to delete the audit configuration for the SVM. |
| 140902425 | Audit configuration is not available for disabling. |
| 140902430 | Audit configuration is not available for enabling. |
| 140902431 | Audit enable failed, audit configuration already enabled for the SVM. |
| 140902432 | Final consolidation is in progress, audit enable failed. |
| 140902445 | Audit disable failed, audit configuration does not exist for the SVM. |
| 140902446 | Audit disable failed, audit configuration does not exist for the SVM. |
| 140902447 | Audit disable failed. |
| 140902456 | The specified log_path does not exist. |
| 140902457 | The log_path must be a directory. |
| 140902458 | The log_path must be a canonical path in the SVM's namespace. |

| Error Code | Description |
| --- | --- |
| 140902459 | The log_path cannot be empty. |
| 140902460 | Rotate size must be greater than or equal to 1024 KB. |
| 140902461 | The destination path must not contain a symbolic link. |
| 140902470 | The log_path exceeds a maximum supported length of characters. |
| 140902471 | The log_path contains an unsupported read-only (DP/LS) volume. |
| 140902472 | The log_path is not a valid destination for the SVM. |
| 140902474 | The log_path contains an unsupported Snaplock volume. |
| 140902478 | The log_path validation failed. |
| 140902478 | The log_path cannot be accessed for validation. |
| 140902490 | Audit configuration is absent for rotate. |
| 140902491 | Failed to rotate audit log. |
| 140902492 | Cannot rotate audit log, auditing is not enabled for this SVM. |

Also see the table of common errors in the Response body overview section of this documentation.
ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 262186 | Field "log.retention.duration" cannot be used with field "log.retention.count" |
| 9699340 | SVM UUID lookup failed |
| 9699407 | Additional fields are provided |

#### Definitions

**See Definitions**

events

| Name | Type | Description |
|---|---|---|
| data | boolean | Data events |
| management | boolean | Management events |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

retention

| Name | Type | Description |
|---|---|---|
| count | integer | Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with "duration". |
| duration | string | Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with "count". |

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

| Name | Type | Description |
|---|---|---|
| days | array[integer] | Specifies the day of the month schedule to rotate audit log. Leave empty for all. |
| hours | array[integer] | Specifies the hourly schedule to rotate audit log. Leave empty for all. |

| Name | Type | Description |
| --- | --- | --- |
| minutes | array[integer] | Specifies the minutes schedule to rotate the audit log. |
| months | array[integer] | Specifies the months schedule to rotate audit log. Leave empty for all. |
| weekdays | array[integer] | Specifies the weekdays schedule to rotate audit log. Leave empty for all. |

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

| Name | Type | Description |
| --- | --- | --- |
| now | boolean | Manually rotates the audit logs. Optional in PATCH only. Not available in POST. |
| schedule | audit_schedule | Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values. |
| size | integer | Rotates logs based on log size in bytes. |

s3_log

| Name | Type | Description |
| --- | --- | --- |
| format | string | Format in which the logs are generated by the consolidation process.<br>Possible values are:<br>• json - ONTAP-specific Json log format.<br>◦ Default value: 1<br>◦ enum: ["json"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |
| retention | retention | |
| rotation | rotation | Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_audit

Auditing for NAS events is a security measure that enables you to track and log certain S3 events on SVMs.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether or not auditing is enabled on the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| events | events | |
| log | s3_log | |
| log_path | string | The audit log destination path where consolidated audit logs are stored. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View CIFS connection information

### View CIFS connection information

#### Overview

ONTAP CIFS connections show functionality is used to display currently established CIFS connections.

### Information on the CIFS connection

- Retrieve the list of the established CIFS connections

#### Example

#### Retrieving established connection information

To retrieve the list of CIFS connections, use the following API.

```
# The API:
GET /protocols/cifs/connections

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/connections?return_timeout=15&return_records=true"
-H "accept: application/json"

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "85d46998-4e5d-11ea-afb1-0050568ec4e4",
        "name": "example_node_name"
      },
      "svm": {
        "uuid": "fc824aa8-4e60-11ea-afb1-0050568ec4e4",
        "name": "vs1"
      },
      "sessions": [
            {
              identifier: 625718873227788312
            }
                  ],
      "identifier": 91842,
      "server_ip": "10.140.70.197",
      "client_ip": "10.74.7.182",
      "client_port": 12345,
      "network_context_id": 3
    },
    {
      "node": {
        "uuid": "85d46998-4e5d-11ea-afb1-0050568ec4e4",
        "name": "example_node_name"
      },
      "svm": {
        "uuid": "fc824aa8-4e60-11ea-afb1-0050568ec4e4",
         "name": "vs1"
      },
      "sessions": [
```

```
            {
              identifier: 625718873227788579
            },
            {
              identifier: 625718873227788577
            }
                    ],
        "identifier": 92080,
        "server_ip": "10.140.70.197",
        "client_ip": "10.140.133.97",
        "client_port": 23413,
        "network_context_id": 5
      }
    ],
    "num_records": 2
  }
```


### Retrieve the CIFS connection information for all SVMs

GET `/protocols/cifs/connections`

**Introduced In:** 9.11

Retrieves the CIFS connection information for all SVMs.

#### Related ONTAP commands

- `vserver cifs connection show`

#### Learn more

- DOC /protocols/cifs/connections

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sessions.identifier | integer | query | False | Filter by sessions.identifier |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| server_ip | string | query | False | Filter by server_ip |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| identifier | integer | query | False | Filter by identifier |
| network_context_id | integer | query | False | Filter by network_context_id |
| client_ip | string | query | False | Filter by client_ip |
| client_port | integer | query | False | Filter by client_port |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_connection] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "client_ip": "10.74.7.182",
      "client_port": 12345,
      "identifier": 22802,
      "network_context_id": 22802,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "server_ip": "10.140.78.248",
      "sessions": [
        {
          "identifier": 4622663542519103507
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

sessions

| Name | Type | Description |
|---|---|---|
| identifier | integer | A unique 64-bit unsigned number represented as string used to represent each SMB session's identifier. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_connection

| Name | Type | Description |
| --- | --- | --- |
| client_ip | string | Specifies IP of the client. |
| client_port | integer | "A unique 32-bit unsigned number used to represent the port number of the connection". |
| identifier | integer | A unique 32-bit unsigned number used to represent each SMB session's connection ID. |
| network_context_id | integer | A unique 32-bit unsigned number used to represent each SMB session's network context ID. |
| node | node | |
| server_ip | string | Specifies the IP address of the SVM. |
| sessions | array[sessions] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View CIFS domain-related information

### View CIFS domain-related information

#### Overview

Displays CIFS domain-related information of all SVMs.

#### Examples

### Retrieving all the fields of CIFS domain configurations of all SVMs

```
# The API:
GET /api/protocols/cifs/domains

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/domains/?fields=*" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "09cc9c9a-c7df-11eb-af15-0050568e403e",
      "name": "vs1"
    },
    "name_mapping": {
      "trusted_domains": [
        "SERVER02.COM",
        "SERVER03.COM"
      ]
    },
    "trust_relationships": [
      {
        "node": {
          "name": "vsNode1",
          "uuid": "a64c0906-c7dd-11eb-af15-0050568e403e"
```

```
        },
        "home_domain": "SERVER02.COM",
        "trusted_domains": [
          "SERVER02.COM"
        ]
      },
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "home_domain": "SERVER02.COM",
        "trusted_domains": [
          "SERVER02.COM"
        ]
      }
    ],
    "discovered_servers": [
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "domain": "server02.com",
        "server_type": "kerberos",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      },
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "domain": "server02.com",
        "server_type": "ms_ldap",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      },
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
```

```
        },
        "domain": "server02.com",
        "server_type": "ms_dc",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      }
    ],
    "preferred_dcs": [
      {
        "fqdn": "server02.com",
        "server_ip": "192.168.20.1"
      }
    ],
    "password_schedule": {
      "schedule_enabled": false,
      "schedule_weekly_interval": 6,
      "schedule_randomized_minute": 120,
      "schedule_description": "Sun@1:00"
    }
  },
  {
    "svm": {
      "uuid": "6dd78167-c907-11eb-b2bf-0050568e7324",
      "name": "vs2"
    },
    "name_mapping": {
      "trusted_domains": [
        "SERVER03.COM",
        "SERVER04.COM"
      ]
    },
    "trust_relationships": [
      {
        "node": {
          "name": "vsNode1",
          "uuid": "a64c0906-c7dd-11eb-af15-0050568e403e"
        },
        "home_domain": "SERVER02.COM",
        "trusted_domains": [
          "SERVER02.COM"
        ]
      },
      {
        "node": {
```

```
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "home_domain": "SERVER02.COM",
        "trusted_domains": [
          "SERVER02.COM"
        ]
      }
    ],
    "discovered_servers": [
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "domain": "server02.com",
        "server_type": "kerberos",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      },
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "domain": "server02.com",
        "server_type": "ms_ldap",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      },
      {
        "node": {
          "name": "vsNode2",
          "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
        },
        "domain": "server02.com",
        "server_type": "ms_dc",
        "server_name": "scspb0659002001",
        "server_ip": "192.168.20.1",
        "preference": "preferred",
        "state": "undetermined"
      }
```

```
      ],
      "preferred_dcs": [
        {
          "fqdn": "server02.com",
          "server_ip": "192.168.20.1"
        }
      ],
      "password_schedule": {
        "schedule_enabled": false,
        "schedule_weekly_interval": 4,
        "schedule_randomized_minute": 120,
        "schedule_description": "Tue@1:00"
      }
    }
  ],
  "num_records": 2
  }
```


---

### Retrieve the CIFS domain-related information of all SVMs

GET `/protocols/cifs/domains`

**Introduced In:** 9.10

Retrieves the CIFS domain-related information of all SVMs.

#### Related ONTAP commands

- `vserver cifs domain preferred-dc show`
- `vserver cifs domain trusts show`
- `vserver cifs domain discovered-servers show`
- `vserver cifs domain name-mapping-search show`
- `vserver cifs domain schedule show`

#### Learn more

- DOC /protocols/cifs/domains

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| password_schedule.schedule_warn_message | string | query | False | Filter by password_schedule.schedule_warn_message |
| password_schedule.schedule_last_changed_time | string | query | False | Filter by password_schedule.schedule_last_changed_time |
| password_schedule.schedule_randomized_minute | integer | query | False | Filter by password_schedule.schedule_randomized_minute<br>• Max value: 180<br>• Min value: 1 |
| password_schedule.schedule_description | string | query | False | Filter by password_schedule.schedule_description |
| password_schedule.schedule_enabled | boolean | query | False | Filter by password_schedule.schedule_enabled |
| password_schedule.schedule_weekly_interval | integer | query | False | Filter by password_schedule.schedule_weekly_interval<br>• Max value: 52<br>• Min value: 1 |
| tenant_id | string | query | False | Filter by tenant_id<br>• Introduced in: 9.16 |
| trust_relationships.node.name | string | query | False | Filter by trust_relationships.node.name |
| trust_relationships.node.uuid | string | query | False | Filter by trust_relationships.node.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| trust_relationships.trusted_domains | string | query | False | Filter by trust_relationships.trusted_domains |
| trust_relationships.home_domain | string | query | False | Filter by trust_relationships.home_domain |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| server_discovery_mode | string | query | False | Filter by server_discovery_mode<br>• Introduced in: 9.13 |
| name_mapping.trusted_domains | string | query | False | Filter by name_mapping.trusted_domains |
| client_id | string | query | False | Filter by client_id<br>• Introduced in: 9.16 |
| preferred_dcs.server_ip | string | query | False | Filter by preferred_dcs.server_ip |
| preferred_dcs.fqdn | string | query | False | Filter by preferred_dcs.fqdn |
| discovered_servers.node.name | string | query | False | Filter by discovered_servers.node.name |
| discovered_servers.node.uuid | string | query | False | Filter by discovered_servers.node.uuid |
| discovered_servers.preference | string | query | False | Filter by discovered_servers.preference |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| discovered_servers.server_name | string | query | False | Filter by discovered_servers.server_name |
| discovered_servers.server_type | string | query | False | Filter by discovered_servers.server_type |
| discovered_servers.server_ip | string | query | False | Filter by discovered_servers.server_ip |
| discovered_servers.domain | string | query | False | Filter by discovered_servers.domain |
| discovered_servers.state | string | query | False | Filter by discovered_servers.state |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of CIFS domain records. |
| records | array[cifs_domain] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
      "discovered_servers": [
        {
          "domain": "test.com",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "preference": "string",
          "server_ip": "string",
          "server_name": "string",
          "server_type": "string",
          "state": "string"
        }
      ],
      "name_mapping": {
        "trusted_domains": [
          "domain.com"
        ]
      },
      "password_schedule": {
        "schedule_day_of_week": "Sunday",
        "schedule_description": "string",
        "schedule_last_changed_time": "string",
        "schedule_time_of_day": 36900,
        "schedule_warn_message": "string"
      },
```

```
      "preferred_dcs": [
        {
          "fqdn": "test.com",
          "server_ip": "4.4.4.4"
        }
      ],
      "server_discovery_mode": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
      "trust_relationships": [
        {
          "home_domain": "string",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "trusted_domains": [
            "string"
          ]
        }
      ]
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

ad_domain

Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account.

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | Fully Qualified Domain Name. |
| node | node | |
| preference | string | Server Preference |
| server_ip | string | Server IP address |
| server_name | string | Server Name |
| server_type | string | Server Type |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Server status |

name_mapping

Specifies the name mapping search records.

| Name | Type | Description |
| --- | --- | --- |
| trusted_domains | array[string] | |

password_schedule

Specifies the password schedule records.

| Name | Type | Description |
| --- | --- | --- |
| schedule_day_of_week | string | Day of the week for password change schedule. |
| schedule_description | string | Schedule description. |
| schedule_enabled | boolean | Is password schedule enabled. |
| schedule_last_changed_time | string | Last successful password change time. |
| schedule_randomized_minute | integer | Minutes within which schedule start can be randomized. |
| schedule_time_of_day | string | Start time for password change schedule. |
| schedule_warn_message | string | Warning message in case job is deleted. |
| schedule_weekly_interval | integer | Interval in weeks for password change schedule. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |

| Name | Type | Description |
| --- | --- | --- |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

trust_relationships

| Name | Type | Description |
| --- | --- | --- |
| home_domain | string | Home Domain Name |
| node | node | |
| trusted_domains | array[string] | Trusted Domain Name |

cifs_domain

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |

| Name | Type | Description |
| --- | --- | --- |
| name_mapping | name_mapping | Specifies the name mapping search records. |
| password_schedule | password_schedule | Specifies the password schedule records. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred DC records. |
| server_discovery_mode | string | Specifies the mode of server discovery. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| trust_relationships | array[trust_relationships] | Specifies the trusted domain records. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage CIFS domain-related information

## Manage CIFS domain-related information

### Overview

Displays CIFS domain-related information of the specified SVM.

### Examples

#### Retrieving all the fields of CIFS domain configurations of a specific SVM

```
# The API:
GET /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/domains/6dd78167-c907-
11eb-b2bf-0050568e7324" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "6dd78167-c907-11eb-b2bf-0050568e7324",
  "name": "vs2"
},
"name_mapping": {
  "trusted_domains": [
    "SERVER03.COM",
    "SERVER04.COM"
  ]
},
"server_discovery_mode": "all",
"trust_relationships": [
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "home_domain": "SERVER02.COM",
    "trusted_domains": [
      "SERVER02.COM"
    ]
  }
],
"discovered_servers": [
  {
```

```
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "kerberos",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  },
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "ms_ldap",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  },
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "ms_dc",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  }
],
"preferred_dcs": [
  {
    "fqdn": "server02.com",
    "server_ip": "192.168.20.1"
  }
],
"password_schedule": {
  "schedule_enabled": false,
  "schedule_weekly_interval": 4,
  "schedule_randomized_minute": 120,
```

```
  "schedule_description": "Tue@1:00"
}
}
```

#### Applying rediscover_trusts query parameter and retrieving all the fields of CIFS domain configurations

```
# The API:
GET /api/protocols/cifs/domains/{svm.uuid}?rediscover_trusts=true

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/domains/6dd78167-c907-
11eb-b2bf-0050568e7324?rediscover_trusts=true" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "6dd78167-c907-11eb-b2bf-0050568e7324",
  "name": "vs2"
},
"name_mapping": {
  "trusted_domains": [
    "SERVER03.COM",
    "SERVER04.COM"
  ]
},
"server_discovery_mode": "all",
"trust_relationships": [
  {
    "node": {
      "name": "vsNode1",
      "uuid": "a64c0906-c7dd-11eb-af15-0050568e403e"
    },
    "home_domain": "SERVER02.COM",
    "trusted_domains": [
      "SERVER02.COM"
    ]
  },
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
```

```
    },
    "home_domain": "SERVER02.COM",
    "trusted_domains": [
      "SERVER02.COM"
    ]
  }
],
"discovered_servers": [
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "kerberos",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  },
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "ms_ldap",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  },
  {
    "node": {
      "name": "vsNode2",
      "uuid": "4d9400f0-c84b-11eb-90ab-0050568e7324"
    },
    "domain": "server02.com",
    "server_type": "ms_dc",
    "server_name": "scspb0659002001",
    "server_ip": "192.168.20.1",
    "preference": "preferred",
    "state": "undetermined"
  }
],
"preferred_dcs": [
```

```
  {
    "fqdn": "server02.com",
    "server_ip": "192.168.20.1"
  }
],
"password_schedule": {
  "schedule_enabled": false,
  "schedule_weekly_interval": 4,
  "schedule_randomized_minute": 120,
  "schedule_description": "Tue@1:00"
}
}
```

#### Modifying the Password Schedule of CIFS domain configurations of a specific SVM

```
# The API:
PATCH /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/domains/a83f816a-c873-
11ed-a3b2-0050568e278e" -H "accept: application/json" -d "{
\"password_schedule\": { \"schedule_enabled\": \"true\" }}"
```

#### Modifying the discovery-mode of CIFS domain configurations of a specific SVM

```
# The API:
PATCH /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/domains/a83f816a-c873-
11ed-a3b2-0050568e278e" -H "accept: application/json" -d "{
\"server_discovery_mode\": \"site\" }"
```

#### Reset the CIFS domain password for a specific SVM

```
# The API:
PATCH /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/domains/a12f816a-c873-
11ed-a3b2-0050568e278e?cifs_password_operation=reset" -H "accept:
application/json" -d "{ \"ad_domain\":{ \"user\": \"administrator\" ,
\"password\": \"admin_password\" }}"
```

#### Reset the CIFS domain password for a specific SVM with hybrid auth user type

```
# The API:
PATCH /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/domains/a12f816a-c873-
11ed-a3b2-0050568e278e?cifs_password_operation=reset" -H "accept:
application/json" -d "{ \"tenant_id\": \"ffffffff-ffff-ffff-ffff-
ffffffffffff\" , \"client_id\": \"ffffffff-zzzz-zzzz-zzzz-ffffffffffff\" ,
\"client_certificate\": \"certificate added to the azure application\" }"
```

#### Change the CIFS domain password for a specific SVM

```
# The API:
PATCH /api/protocols/cifs/domains/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/domains/a12f816a-c873-
11ed-a3b2-0050568e278e?cifs_password_operation=reset" -H "accept:
application/json" -d "{ }"
```

#### Retrieve the CIFS domain-related information for an SVM

GET `/protocols/cifs/domains/{svm.uuid}`

**Introduced In:** 9.10

Retrieves the CIFS domain-related information of the specified SVM.

#### Important notes

GET operation with query parameter `rediscover_trusts` and `reset_discovered_servers` returns available CIFS domain configurations and also triggers trusts rediscovery and discovered servers reset asynchronously for that SVM.

#### Related ONTAP commands

- `vserver cifs domain preferred-dc show`
- `vserver cifs domain trusts show`
- `vserver cifs domain discovered-servers show`
- `vserver cifs domain name-mapping-search show`
- `vserver cifs domain password schedule show`
- `vserver cifs domain discovered-servers discovery-mode show`

#### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rediscover_trusts | boolean | query | False | Force the discovery of trusted domains.<br>• Default value: |
| reset_discovered_servers | boolean | query | False | Force a rediscovery.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| name_mapping | name_mapping | Specifies the name mapping search records. |
| password_schedule | password_schedule | Specifies the password schedule records. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred DC records. |
| server_discovery_mode | string | Specifies the mode of server discovery. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| trust_relationships | array[trust_relationships] | Specifies the trusted domain records. |

**Example response**


```
{
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "discovered_servers": [
    {
      "domain": "test.com",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "string",
      "server_ip": "string",
      "server_name": "string",
      "server_type": "string",
      "state": "string"
    }
  ],
  "name_mapping": {
    "trusted_domains": [
      "domain.com"
    ]
  },
  "password_schedule": {
    "schedule_day_of_week": "Sunday",
    "schedule_description": "string",
    "schedule_last_changed_time": "string",
    "schedule_time_of_day": 36900,
    "schedule_warn_message": "string"
  },
  "preferred_dcs": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ],
  "server_discovery_mode": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
    "trust_relationships": [
      {
        "home_domain": "string",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "trusted_domains": [
          "string"
        ]
      }
    ]
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 656463 | SVM UUID must be provided for a query on the field rediscover_trusts and reset_discovered_servers. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

ad_domain

Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account.

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | Fully Qualified Domain Name. |
| node | node | |
| preference | string | Server Preference |
| server_ip | string | Server IP address |
| server_name | string | Server Name |
| server_type | string | Server Type |
| state | string | Server status |

name_mapping

Specifies the name mapping search records.

| Name | Type | Description |
| --- | --- | --- |
| trusted_domains | array[string] | |

password_schedule

Specifies the password schedule records.

| Name | Type | Description |
| --- | --- | --- |
| schedule_day_of_week | string | Day of the week for password change schedule. |
| schedule_description | string | Schedule description. |
| schedule_enabled | boolean | Is password schedule enabled. |
| schedule_last_changed_time | string | Last successful password change time. |
| schedule_randomized_minute | integer | Minutes within which schedule start can be randomized. |
| schedule_time_of_day | string | Start time for password change schedule. |
| schedule_warn_message | string | Warning message in case job is deleted. |
| schedule_weekly_interval | integer | Interval in weeks for password change schedule. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

trust_relationships

| Name | Type | Description |
| --- | --- | --- |
| home_domain | string | Home Domain Name |
| node | node | |
| trusted_domains | array[string] | Trusted Domain Name |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Modify the CIFS domain-related information for an SVM**

PATCH `/protocols/cifs/domains/{svm.uuid}`

**Introduced In:** 9.13

Modifies the CIFS domain-related information of the specified SVM.

#### Related ONTAP commands

- `vserver cifs domain password schedule modify`
- `vserver cifs domain discovered-servers discovery-mode modify`
- `vserver cifs domain password reset`
- `vserver cifs domain password change`

#### Important notes

- If the query is set to cifs_password_operation=reset and ad_domain_user and ad_domain_password is included in the body,a CIFS password reset is executed.
- If the query is set to cifs_password_operation=reset and tenant_id, client_id and client_certificate is included in the body,a CIFS password reset is executed in Azure EntraId for hybrid user.
- If the body is empty when the query is set to cifs_password_operation=reset, a CIFS password change is executed.

#### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cifs_password_operation | string | query | False | Specifies the operation to be performed on the domain password.<br>• enum: ["reset"]<br>• Introduced in: 9.14 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account. |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| name_mapping | name_mapping | Specifies the name mapping search records. |
| password_schedule | password_schedule | Specifies the password schedule records. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred DC records. |
| server_discovery_mode | string | Specifies the mode of server discovery. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| trust_relationships | array[trust_relationships] | Specifies the trusted domain records. |

**Example request**

```
{
  "ad_domain": {
    "password": "string",
    "user": "string"
  },
  "client_certificate": "PEM Cert",
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "discovered_servers": [
    {
      "domain": "test.com",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "preference": "string",
      "server_ip": "string",
      "server_name": "string",
      "server_type": "string",
      "state": "string"
    }
  ],
  "name_mapping": {
    "trusted_domains": [
      "domain.com"
    ]
  },
  "password_schedule": {
    "schedule_day_of_week": "Sunday",
    "schedule_description": "string",
    "schedule_last_changed_time": "string",
    "schedule_time_of_day": 36900,
    "schedule_warn_message": "string"
  },
  "preferred_dcs": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4"
    }
  ],
  "server_discovery_mode": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
```

```
  "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
  "trust_relationships": [
    {
      "home_domain": "string",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "trusted_domains": [
        "string"
      ]
    }
  ]
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262278 | Domain user and password are required fields. |
| 655513 | Password update failed due to invalid credentials. |
| 656392 | Preferred DC's need to be configured if CIFS discovery mode is set to none. |
| 656393 | Default site needs to be configured if CIFS discovery mode is set to site. |
| 656487 | Missing fields when doing CIFS operation with hybrid auth user type. |

**Definitions**

**See Definitions**

ad_domain

Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account.

| Name | Type | Description |
|---|---|---|
| password | string | The account password of the user with permissions to reset the password in the organizational unit for the machine account. |
| user | string | The username of a user with permissions to reset the password in the organizational unit for the machine account. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

discovered_servers

| Name | Type | Description |
|---|---|---|
| domain | string | Fully Qualified Domain Name. |
| node | node | |
| preference | string | Server Preference |
| server_ip | string | Server IP address |
| server_name | string | Server Name |
| server_type | string | Server Type |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Server status |

name_mapping

Specifies the name mapping search records.

| Name | Type | Description |
| --- | --- | --- |
| trusted_domains | array[string] | |

password_schedule

Specifies the password schedule records.

| Name | Type | Description |
| --- | --- | --- |
| schedule_day_of_week | string | Day of the week for password change schedule. |
| schedule_description | string | Schedule description. |
| schedule_enabled | boolean | Is password schedule enabled. |
| schedule_last_changed_time | string | Last successful password change time. |
| schedule_randomized_minute | integer | Minutes within which schedule start can be randomized. |
| schedule_time_of_day | string | Start time for password change schedule. |
| schedule_warn_message | string | Warning message in case job is deleted. |
| schedule_weekly_interval | integer | Interval in weeks for password change schedule. |

preferred_dcs

| Name | Type | Description |
| --- | --- | --- |
| fqdn | string | Fully Qualified Domain Name. |

| Name | Type | Description |
| --- | --- | --- |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

trust_relationships

| Name | Type | Description |
| --- | --- | --- |
| home_domain | string | Home Domain Name |
| node | node | |
| trusted_domains | array[string] | Trusted Domain Name |

cifs_domain

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | Specifies the credentials of a user with permissions to reset the password in the organizational unit for the machine account. |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |

| Name | Type | Description |
| --- | --- | --- |
| discovered_servers | array[discovered_servers] | Specifies the discovered servers records. |
| name_mapping | name_mapping | Specifies the name mapping search records. |
| password_schedule | password_schedule | Specifies the password schedule records. |
| preferred_dcs | array[preferred_dcs] | Specifies the preferred DC records. |
| server_discovery_mode | string | Specifies the mode of server discovery. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| trust_relationships | array[trust_relationships] | Specifies the trusted domain records. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage CIFS DC configuration

### Manage CIFS DC configuration

#### Overview

You can use this API to display a CIFS domain preferred DC configuration of an SVM.

#### Retrieving all CIFS domain preferred DC configurations of an SVM

---

The CIFS domain preferred DC GET endpoint retrieves all the configurations for a specific SVM.

#### Examples

### Retrieving all the fields of all CIFS domain preferred DC configurations of an SVM

---

```
# The API:
GET /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-abc9-
11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fields=*" -H "accept:
application/json"

# The response:
{
 "records": [
{
  "svm": {
    "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
    "name": "svm1"
  },
  "fqdn": "host1",
  "server_ip": "4.4.4.4"
},
{
  "svm": {
    "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
    "name": "svm1"
  },
  "fqdn": "host2",
  "server_ip": "11.11.11.11"
}
],
 "num_records": 2
}
```

#### Retrieving the CIFS domain preferred DC configuration of a specific SVM, "fqdn" and "server_ip"

```
# The API:
GET /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-
controllers/{fqdn}/{server_ip}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-abc9-
11eb-8de3-0050568eb0c4/preferred-domain-controllers/host1/4.4.4.4" -H
"accept: application/json"

# The response:
{
"svm": {
  "uuid": "1226670c-abc9-11eb-8de3-0050568eb0c4",
  "name": "svm1"
},
"fqdn": "host1",
"server_ip": "4.4.4.4"
}
```

---

#### Creating a new CIFS domain preferred DC configuration

The CIFS domain preferred DC POST endpoint creates a new configuration. Both bulk and instance POST is supported.

#### Examples

### Creating a single CIFS domain preferred DC configuration

---

```
# The API:
POST /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-abc9-
11eb-8de3-0050568eb0c4/preferred-domain-
controllers?skip_config_validation=true&return_records=false" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fqdn\":
\"testing.com\", \"server_ip\": \"1.1.1.1\"}"
```

---

#### Deleting an existing CIFS domain preferred DC configuration

The CIFS domain preferred DC DELETE endpoint deletes an existing configuration. Both bulk and instance delete is supported.

#### Examples

### Deleting the CIFS domain preferred DC configuration of a specific SVM, "fqdn" and "server_ip"

```
# The API:
DELETE /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-
controllers/{fqdn}/{server_ip}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/sample/4.4.4.4"
```

```
### Deleting the CIFS domain preferred DC configurations of a specific SVM
and "fqdn"
```

```
# The API:
DELETE /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fqdn=sampl"
```

### Deleting all CIFS domain preferred DC configurations of a specific SVM

```
# The API:
DELETE /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-controllers/?fqdn=*"
```

#### Deleting the CIFS domain preferred DC configurations of a specific SVM, "fqdn" and set of "server_ips"

```
# The API:
DELETE /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-
controllers/?fqdn=sample&server_ip=3.3.3.3&#124;4.4.4.4&#124;1.1.1.1&#124;
2.2.2.2"
```

#### Deleting the CIFS domain preferred DC configurations of a specific SVM and set of "server_ips"

```
# The API:
DELETE /api/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/domains/1226670c-
abc9-11eb-8de3-0050568eb0c4/preferred-domain-
controllers/?server_ip=3.3.3.3&#124;4.4.4.4&#124;1.1.1.1&#124;2.2.2.2"
```

#### Retrieve the CIFS domain preferred DC configuration for an SVM

GET `/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers`

**Introduced In:** 9.10

Retrieves the CIFS domain preferred DC configuration of an SVM.

##### Related ONTAP commands

- `vserver cifs domain preferred-dc show`
- `vserver cifs domain preferred-dc check`

##### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| need_status | boolean | query | False | Retrieves the status of the preferred DCs.<br>• Default value:<br>• Introduced in: 9.12 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.16 |
| server_ip | string | query | False | Filter by server_ip |
| status.details | string | query | False | Filter by status.details<br>• Introduced in: 9.12 |
| status.reachable | boolean | query | False | Filter by status.reachable<br>• Introduced in: 9.12 |
| fqdn | string | query | False | Filter by fqdn |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of CIFS preferred domain controller records. |
| records | array[cifs_domain_preferred_dc] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "fqdn": "test.com",
      "server_ip": "4.4.4.4",
      "status": {
        "details": "Response time (msec): 111",
        "reachable": 1
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

status

Status of CIFS preferred domain controller.

| Name | Type | Description |
|---|---|---|
| details | string | Provides a detailed description of the state if the state is 'down' or the response time of the DNS server if the state is 'up'. |
| reachable | boolean | Indicates whether or not the domain controller is reachable. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_domain_preferred_dc

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |
| status | status | Status of CIFS preferred domain controller. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the CIFS domain preferred DC configuration for an SVM

POST `/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers`

**Introduced In:** 9.10

Creates a CIFS preferred DC configuration for an SVM.

#### Important notes

- In the case of bulk POST requests, the create operation should be performed serially since there can be interdependence between records. In order to avoid issues, it is advisable to always use the query

parameter "serial_records=true".

#### Required properties

- `svm.uuid` - Existing SVM in which to create the preferred-dc.
- `domain` - Fully Qualified Domain Name.
- `server_ip` - IPv4/IPv6 address of the Preferred Domain Controller.

#### The following parameters are optional:

- skip_config_validation

#### Related ONTAP commands

- `vserver cifs domain preferred-dc add`

#### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| skip_config_validation | boolean | query | False | Skip the validation of the specified preferred DC configuration.<br>• Default value: |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |

| Name | Type | Description |
|---|---|---|
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |
| status | status | Status of CIFS preferred domain controller. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4",
  "status": {
    "details": "Response time (msec): 111",
    "reachable": 1
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |
| status | status | Status of CIFS preferred domain controller. |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4",
  "status": {
    "details": "Response time (msec): 111",
    "reachable": 1
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621516 | Only data SVMs allowed. |
| 655918 | The fully qualified domain name cannot be longer than 254 bytes. |
| 656408 | RPC failure occurred during the CIFS preferred-dc configuration validation. |
| 656407 | Failed to validate CIFS preferred-dc for domain. Reason: Configuration not found at SecD. Contact technical support for assistance. |
| 655366 | Invalid domain controller. |

| Error Code | Description |
| --- | --- |
| 655506 | Failed to add preferred-dc. |

**Definitions**

**See Definitions**

status

Status of CIFS preferred domain controller.

| Name | Type | Description |
|---|---|---|
| details | string | Provides a detailed description of the state if the state is 'down' or the response time of the DNS server if the state is 'up'. |
| reachable | boolean | Indicates whether or not the domain controller is reachable. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_domain_preferred_dc

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |

| Name | Type | Description |
|---|---|---|
| status | status | Status of CIFS preferred domain controller. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the CIFS domain preferred DC configuration for an SVM and domain

DELETE `/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers/{fqdn}/{server_ip}`

**Introduced In:** 9.10

Deletes the CIFS domain preferred DC configuration of the specified SVM and domain.

#### Related ONTAP commands

- `vserver cifs domain preferred-dc delete`

#### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fqdn | string | path | True | Fully Qualified Domain Name |
| server_ip | string | path | True | Domain Controller IP address |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655507 | Failed to remove preferred-dc. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the CIFS domain preferred DC configuration for an SVM and domain

GET `/protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers/{fqdn}/{server_ip}`

**Introduced In:** 9.10

Retrieves the CIFS domain preferred DC configuration of an SVM.

#### Related ONTAP commands

- `vserver cifs domain preferred-dc show`
- `vserver cifs domain preferred-dc check`

#### Learn more

- DOC /protocols/cifs/domains/{svm.uuid}/preferred-domain-controllers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fqdn | string | path | True | Fully Qualified Domain Name |
| server_ip | string | path | True | Domain Controller IP address |
| need_status | boolean | query | False | Retrieves the status of the specified preferred DC.<br>• Introduced in: 9.12<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| fqdn | string | Fully Qualified Domain Name. |

| Name | Type | Description |
|---|---|---|
| server_ip | string | IP address of the preferred domain controller (DC). The address can be either an IPv4 or an IPv6 address. |
| status | status | Status of CIFS preferred domain controller. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "fqdn": "test.com",
  "server_ip": "4.4.4.4",
  "status": {
    "details": "Response time (msec): 111",
    "reachable": 1
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

status

Status of CIFS preferred domain controller.

| Name | Type | Description |
|---|---|---|
| details | string | Provides a detailed description of the state if the state is 'down' or the response time of the DNS server if the state is 'up'. |
| reachable | boolean | Indicates whether or not the domain controller is reachable. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve unapplied group policy objects for all SVMs

GET `/protocols/cifs/group-policies`

**Introduced In:** 9.12

Retrieves group policy objects that are yet to be applied for all SVMs.

#### Related ONTAP commands

- `vserver cifs group-policy show-defined`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| registry_settings.branchcache.supported_hash_version | string | query | False | Filter by registry_settings.branchcache.supported_hash_version |
| registry_settings.branchcache.hash_publication_mode | string | query | False | Filter by registry_settings.branchcache.hash_publication_mode |
| registry_settings.refresh_time_random_offset | string | query | False | Filter by registry_settings.refresh_time_random_offset |
| registry_settings.refresh_time_interval | string | query | False | Filter by registry_settings.refresh_time_interval |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• minLength: 1 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| central_access_policy_settings | string | query | False | Filter by central_access_policy_settings |
| file_system_path | string | query | False | Filter by file_system_path |
| enabled | boolean | query | False | Filter by enabled |
| central_access_policy_staging_audit_type | string | query | False | Filter by central_access_policy_staging_audit_type |
| ldap_path | string | query | False | Filter by ldap_path |
| extensions | string | query | False | Filter by extensions |
| security_settings.event_audit_settings.object_access_type | string | query | False | Filter by security_settings.event_audit_settings.object_access_type |
| security_settings.event_audit_settings.logon_type | string | query | False | Filter by security_settings.event_audit_settings.logon_type |
| security_settings.restrict_anonymous.no_enumeration_of_sam_accounts | boolean | query | False | Filter by security_settings.restrict_anonymous.no_enumeration_of_sam_accounts |
| security_settings.restrict_anonymous.combined_restriction_for_anonymous_user | string | query | False | Filter by security_settings.restrict_anonymous.combined_restriction_for_anonymous_user |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_settings.restrict_anonymous.no_enumeration_of_sam_accounts_and_shares | boolean | query | False | Filter by security_settings.restrict_anonymous.no_enumeration_of_sam_accounts_and_shares |
| security_settings.restrict_anonymous.anonymous_access_to_shares_and_named_pipes_restricted | boolean | query | False | Filter by security_settings.restrict_anonymous.anonymous_access_to_shares_and_named_pipes_restricted |
| security_settings.files_or_folders | string | query | False | Filter by security_settings.files_or_folders |
| security_settings.registry_values.signing_required | boolean | query | False | Filter by security_settings.registry_values.signing_required |
| security_settings.privilege_rights.change_notify_users | string | query | False | Filter by security_settings.privilege_rights.change_notify_users |
| security_settings.privilege_rights.security_privilege_users | string | query | False | Filter by security_settings.privilege_rights.security_privilege_users |
| security_settings.privilege_rights.take_ownership_users | string | query | False | Filter by security_settings.privilege_rights.take_ownership_users |
| security_settings.restricted_groups | string | query | False | Filter by security_settings.restricted_groups |
| security_settings.kerberos.max_ticket_age | string | query | False | Filter by security_settings.kerberos.max_ticket_age |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_settings.kerberos.max_renew_age | string | query | False | Filter by security_settings.kerberos.max_renew_age |
| security_settings.kerberos.max_clock_skew | string | query | False | Filter by security_settings.kerberos.max_clock_skew |
| security_settings.event_log_settings.max_size | integer | query | False | Filter by security_settings.event_log_settings.max_size |
| security_settings.event_log_settings.retention_method | string | query | False | Filter by security_settings.event_log_settings.retention_method |
| link | string | query | False | Filter by link |
| version | integer | query | False | Filter by version |
| index | integer | query | False | Filter by index |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of central access rules. |
| records | array[policies_and_rules_to_be_applied] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "to_be_applied": {
        "access_policies": [
          {
            "create_time": "2018-01-01 12:00:00 -0400",
            "description": "policy #1",
            "member_rules": [
              "r1",
              "r2"
            ],
            "name": "p1",
            "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
            "svm": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "update_time": "2018-01-01 12:00:00 -0400"
          }
        ],
```

```
        "access_rules": [
          {
            "create_time": "2018-01-01 12:00:00 -0400",
            "current_permission": "O:SYG:SYD:AR(A;;FA;;;WD)",
            "description": "rule #1",
            "name": "p1",
            "proposed_permission":
"O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)",
            "resource_criteria": "department",
            "svm": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "update_time": "2018-01-01 12:00:00 -0400"
          }
        ],
        "objects": [
          {
            "central_access_policy_settings": [
              "p1",
              "p2"
            ],
            "central_access_policy_staging_audit_type": "none",
            "extensions": [
              "audit",
              "security"
            ],
            "file_system_path": "\\test.com\\SysVol\\test.com\\
policies\\{42474212-3f9d-4489-ae01-6fcf4f805d4c}",
            "index": 1,
            "ldap_path": "cn={42474212-3f9d-4489-ae01-
6fcf4f805d4c},cn=policies,cn=system,DC=TEST,DC=COM",
            "link": "domain",
            "name": "test_policy",
            "registry_settings": {
              "branchcache": {
                "hash_publication_mode": "disabled",
                "supported_hash_version": "version1"
              },
              "refresh_time_interval": "P15M",
              "refresh_time_random_offset": "P1D"
```

```
            },
            "security_settings": {
              "event_audit_settings": {
                "logon_type": "failure",
                "object_access_type": "failure"
              },
              "event_log_settings": {
                "max_size": 2048,
                "retention_method": "do_not_overwrite"
              },
              "files_or_folders": [
                "/vol1/home",
                "/vol1/dir1"
              ],
              "kerberos": {
                "max_clock_skew": "P15M",
                "max_renew_age": "P2D",
                "max_ticket_age": "P24H"
              },
              "privilege_rights": {
                "change_notify_users": [
                  "usr1",
                  "usr2"
                ],
                "security_privilege_users": [
                  "usr1",
                  "usr2"
                ],
                "take_ownership_users": [
                  "usr1",
                  "usr2"
                ]
              },
              "restrict_anonymous": {
                "combined_restriction_for_anonymous_user": "no_access"
              },
              "restricted_groups": [
                "test_grp1",
                "test_grp2"
              ]
            },
            "svm": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
```

```
                },
                "name": "svm1",
                "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
              },
              "uuid": "42474212-3f9d-4489-ae01-6fcf4f805d4c",
              "version": 7
            }
          ],
          "restricted_groups": [
            {
              "group_name": "test_group",
              "link": "domain",
              "members": [
                "DOMAIN/test_user",
                "DOMAIN/user2"
              ],
              "memberships": [
                "DOMAIN/AdministratorGrp",
                "DOMAIN/deptMark"
              ],
              "policy_name": "test_policy",
              "svm": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "svm1",
                "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
              },
              "version": 7
            }
          ]
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object_central_access_policy

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| description | string | Description about the policy. |
| member_rules | array[string] | Names of all central access rules applied to members. |
| name | string | |
| sid | string | Security ID, unique identifier of the central policy. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

group_policy_object_central_access_rule

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| current_permission | string | Effective security policy in security descriptor definition language format. |
| description | string | Description about the policy. |
| name | string | |
| proposed_permission | string | Proposed security policy in security descriptor definition language format. |
| resource_criteria | string | Criteria to scope resources for which access rules apply. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

group_policy_object_branchcache

| Name | Type | Description |
|---|---|---|
| hash_publication_mode | string | Hash publication mode. |
| supported_hash_version | string | Hash version. |

group_policy_object_registry_setting

| Name | Type | Description |
|---|---|---|
| branchcache | group_policy_object_branchcache | |
| refresh_time_interval | string | Refresh time interval in ISO-8601 format. |
| refresh_time_random_offset | string | Random offset in ISO-8601 format. |

group_policy_object_event_audit

| Name | Type | Description |
|---|---|---|
| logon_type | string | Type of logon event to be audited. |
| object_access_type | string | Type of object access to be audited. |

group_policy_object_event_log

| Name | Type | Description |
|---|---|---|
| max_size | integer | Maximum size of security log, in kilobytes. |
| retention_method | string | Audit log retention method. |

group_policy_object_kerberos

| Name | Type | Description |
|---|---|---|
| max_clock_skew | string | Kerberos clock skew in ISO-8601 format. |
| max_renew_age | string | Kerberos max renew age in ISO-8601 format. |

| Name | Type | Description |
| --- | --- | --- |
| max_ticket_age | string | Kerberos max ticket age in ISO-8601 format. |

group_policy_object_privilege_right

| Name | Type | Description |
| --- | --- | --- |
| change_notify_users | array[string] | Users with traversing bypass privileges. |
| security_privilege_users | array[string] | Users with security privileges. |
| take_ownership_users | array[string] | Users who can take ownership of securable objects. |

group_policy_object_registry_value

| Name | Type | Description |
| --- | --- | --- |
| signing_required | boolean | SMB signing required. |

group_policy_object_restrict_anonymous

| Name | Type | Description |
| --- | --- | --- |
| anonymous_access_to_shares_and_named_pipes_restricted | boolean | Restrict anonymous access to shares and named pipes. |
| combined_restriction_for_anonymous_user | string | Combined restriction for anonymous user. |
| no_enumeration_of_sam_accounts | boolean | No enumeration of SAM accounts. |
| no_enumeration_of_sam_accounts_and_shares | boolean | No enumeration of SAM accounts and shares. |

group_policy_object_security_setting

| Name | Type | Description |
| --- | --- | --- |
| event_audit_settings | group_policy_object_event_audit | |
| event_log_settings | group_policy_object_event_log | |
| files_or_folders | array[string] | Files/Directories for file security. |

| Name | Type | Description |
|---|---|---|
| kerberos | group_policy_object_kerberos | |
| privilege_rights | group_policy_object_privilege_right | |
| registry_values | group_policy_object_registry_value | |
| restrict_anonymous | group_policy_object_restrict_anonymous | |
| restricted_groups | array[string] | List of restricted groups. |

group_policy_object

| Name | Type | Description |
|---|---|---|
| central_access_policy_settings | array[string] | List of central access policies. |
| central_access_policy_staging_audit_type | string | Types of events to be audited. |
| enabled | boolean | Specifies whether group policies are enabled for the SVM. |
| extensions | array[string] | List of extensions. |
| file_system_path | string | File system path. |
| index | integer | Group policy object index. |
| ldap_path | string | LDAP path to the GPO. |
| link | string | Link info. |
| name | string | |
| registry_settings | group_policy_object_registry_setting | |
| security_settings | group_policy_object_security_setting | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| uuid | string | Policy UUID. |
| version | integer | Group policy object version. |

group_policy_object_restricted_group

| Name | Type | Description |
|---|---|---|
| group_name | string | |
| link | string | Link info. |
| members | array[string] | Members of the group. |
| memberships | array[string] | Group is member of Group/OU. |
| policy_name | string | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| version | integer | Group policy object version. |

to_be_applied

| Name | Type | Description |
|---|---|---|
| access_policies | array[group_policy_object_central_access_policy] | |
| access_rules | array[group_policy_object_central_access_rule] | |
| objects | array[group_policy_object] | |
| restricted_groups | array[group_policy_object_restricted_group] | |

policies_and_rules_to_be_applied

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| to_be_applied | to_be_applied | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve unapplied group policy objects for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}`

**Introduced In:** 9.12

Retrieves group policy objects that are yet to be applied for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy show-defined`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| to_be_applied | to_be_applied | |

**Example response**


```
{
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "to_be_applied": {
    "access_policies": [
      {
        "create_time": "2018-01-01 12:00:00 -0400",
        "description": "policy #1",
        "member_rules": [
          "r1",
          "r2"
        ],
        "name": "p1",
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "update_time": "2018-01-01 12:00:00 -0400"
      }
    ],
    "access_rules": [
      {
        "create_time": "2018-01-01 12:00:00 -0400",
        "current_permission": "O:SYG:SYD:AR(A;;FA;;;WD)",
        "description": "rule #1",
        "name": "p1",
        "proposed_permission":
"O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)",
        "resource_criteria": "department",
        "svm": {
          "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "update_time": "2018-01-01 12:00:00 -0400"
      }
    ],
    "objects": [
      {
        "central_access_policy_settings": [
          "p1",
          "p2"
        ],
        "central_access_policy_staging_audit_type": "none",
        "extensions": [
          "audit",
          "security"
        ],
        "file_system_path": "\\test.com\\SysVol\\test.com\\policies
\\{42474212-3f9d-4489-ae01-6fcf4f805d4c}",
        "index": 1,
        "ldap_path": "cn={42474212-3f9d-4489-ae01-
6fcf4f805d4c},cn=policies,cn=system,DC=TEST,DC=COM",
        "link": "domain",
        "name": "test_policy",
        "registry_settings": {
          "branchcache": {
            "hash_publication_mode": "disabled",
            "supported_hash_version": "version1"
          },
          "refresh_time_interval": "P15M",
          "refresh_time_random_offset": "P1D"
        },
        "security_settings": {
          "event_audit_settings": {
            "logon_type": "failure",
            "object_access_type": "failure"
          },
          "event_log_settings": {
            "max_size": 2048,
            "retention_method": "do_not_overwrite"
          },
          "files_or_folders": [
```

```
            "/vol1/home",
            "/vol1/dir1"
          ],
          "kerberos": {
            "max_clock_skew": "P15M",
            "max_renew_age": "P2D",
            "max_ticket_age": "P24H"
          },
          "privilege_rights": {
            "change_notify_users": [
              "usr1",
              "usr2"
            ],
            "security_privilege_users": [
              "usr1",
              "usr2"
            ],
            "take_ownership_users": [
              "usr1",
              "usr2"
            ]
          },
          "restrict_anonymous": {
            "combined_restriction_for_anonymous_user": "no_access"
          },
          "restricted_groups": [
            "test_grp1",
            "test_grp2"
          ]
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "42474212-3f9d-4489-ae01-6fcf4f805d4c",
        "version": 7
      }
    ],
    "restricted_groups": [
      {
        "group_name": "test_group",
```

```
        "link": "domain",
        "members": [
          "DOMAIN/test_user",
          "DOMAIN/user2"
        ],
        "memberships": [
          "DOMAIN/AdministratorGrp",
          "DOMAIN/deptMark"
        ],
        "policy_name": "test_policy",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "version": 7
      }
    ]
  }
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object_central_access_policy

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |

| Name | Type | Description |
| --- | --- | --- |
| description | string | Description about the policy. |
| member_rules | array[string] | Names of all central access rules applied to members. |
| name | string | |
| sid | string | Security ID, unique identifier of the central policy. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

group_policy_object_central_access_rule

| Name | Type | Description |
| --- | --- | --- |
| create_time | string | Policy creation timestamp. |
| current_permission | string | Effective security policy in security descriptor definition language format. |
| description | string | Description about the policy. |
| name | string | |
| proposed_permission | string | Proposed security policy in security descriptor definition language format. |
| resource_criteria | string | Criteria to scope resources for which access rules apply. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

group_policy_object_branchcache

| Name | Type | Description |
| --- | --- | --- |
| hash_publication_mode | string | Hash publication mode. |
| supported_hash_version | string | Hash version. |

group_policy_object_registry_setting

| Name | Type | Description |
| --- | --- | --- |
| branchcache | group_policy_object_branchcache | |
| refresh_time_interval | string | Refresh time interval in ISO-8601 format. |
| refresh_time_random_offset | string | Random offset in ISO-8601 format. |

group_policy_object_event_audit

| Name | Type | Description |
| --- | --- | --- |
| logon_type | string | Type of logon event to be audited. |
| object_access_type | string | Type of object access to be audited. |

group_policy_object_event_log

| Name | Type | Description |
| --- | --- | --- |
| max_size | integer | Maximum size of security log, in kilobytes. |
| retention_method | string | Audit log retention method. |

group_policy_object_kerberos

| Name | Type | Description |
| --- | --- | --- |
| max_clock_skew | string | Kerberos clock skew in ISO-8601 format. |
| max_renew_age | string | Kerberos max renew age in ISO-8601 format. |

| Name | Type | Description |
| --- | --- | --- |
| max_ticket_age | string | Kerberos max ticket age in ISO-8601 format. |

group_policy_object_privilege_right

| Name | Type | Description |
| --- | --- | --- |
| change_notify_users | array[string] | Users with traversing bypass privileges. |
| security_privilege_users | array[string] | Users with security privileges. |
| take_ownership_users | array[string] | Users who can take ownership of securable objects. |

group_policy_object_registry_value

| Name | Type | Description |
| --- | --- | --- |
| signing_required | boolean | SMB signing required. |

group_policy_object_restrict_anonymous

| Name | Type | Description |
| --- | --- | --- |
| anonymous_access_to_shares_and_named_pipes_restricted | boolean | Restrict anonymous access to shares and named pipes. |
| combined_restriction_for_anonymous_user | string | Combined restriction for anonymous user. |
| no_enumeration_of_sam_accounts | boolean | No enumeration of SAM accounts. |
| no_enumeration_of_sam_accounts_and_shares | boolean | No enumeration of SAM accounts and shares. |

group_policy_object_security_setting

| Name | Type | Description |
| --- | --- | --- |
| event_audit_settings | group_policy_object_event_audit | |
| event_log_settings | group_policy_object_event_log | |
| files_or_folders | array[string] | Files/Directories for file security. |

| Name | Type | Description |
|---|---|---|
| kerberos | group_policy_object_kerberos | |
| privilege_rights | group_policy_object_privilege_right | |
| registry_values | group_policy_object_registry_value | |
| restrict_anonymous | group_policy_object_restrict_anonymous | |
| restricted_groups | array[string] | List of restricted groups. |

group_policy_object

| Name | Type | Description |
|---|---|---|
| central_access_policy_settings | array[string] | List of central access policies. |
| central_access_policy_staging_audit_type | string | Types of events to be audited. |
| enabled | boolean | Specifies whether group policies are enabled for the SVM. |
| extensions | array[string] | List of extensions. |
| file_system_path | string | File system path. |
| index | integer | Group policy object index. |
| ldap_path | string | LDAP path to the GPO. |
| link | string | Link info. |
| name | string | |
| registry_settings | group_policy_object_registry_setting | |
| security_settings | group_policy_object_security_setting | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| uuid | string | Policy UUID. |
| version | integer | Group policy object version. |

group_policy_object_restricted_group

| Name | Type | Description |
|---|---|---|
| group_name | string | |
| link | string | Link info. |
| members | array[string] | Members of the group. |
| memberships | array[string] | Group is member of Group/OU. |
| policy_name | string | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| version | integer | Group policy object version. |

to_be_applied

| Name | Type | Description |
|---|---|---|
| access_policies | array[group_policy_object_central_access_policy] | |
| access_rules | array[group_policy_object_central_access_rule] | |
| objects | array[group_policy_object] | |
| restricted_groups | array[group_policy_object_restricted_group] | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
|---|---|---|
| target | string | The target parameter that caused the error. |

### Create a background task to update the GPO settings for an SVM

PATCH `/protocols/cifs/group-policies/{svm.uuid}`

**Introduced In:** 9.12

Creates a background task to update the GPO settings for the specified SVM.
Note: The group policy can be enabled or disabled using "group_policy_object_enabled" field in PATCH "/protocols/cifs/services/{svm.uuid}" API.

#### Related ONTAP commands

- `vserver cifs group-policy update`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| apply_updates | boolean | query | False | Indicates whether to apply updates or not.<br>• Default value: 1 |
| force_reapply_all_settings | boolean | query | False | Indicates whether or not a re-applying of all settings is enforced.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4456649 | Unable to get Group Policy information for SVM. |
| 4456650 | Group Policy is disabled for SVM. |
| 4456651 | Unable to get CIFS server information for SVM. |
| 4456652 | CIFS server is down for SVM. |
| 4456851 | An update is already in progress for specified SVM. Wait a few minutes, then try the command again. |
| 4456852 | Failed to set forced update for specified SVM. Reason: Internal error. Starting normal update. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve applied central access policies for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/central-access-policies`

**Introduced In:** 9.12

Retrieves applied central access policies for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy central-access-policy show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• minLength: 1 |
| update_time | string | query | False | Filter by update_time |
| svm.name | string | query | False | Filter by svm.name |
| description | string | query | False | Filter by description |
| create_time | string | query | False | Filter by create_time |
| sid | string | query | False | Filter by sid |
| member_rules | string | query | False | Filter by member_rules |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of central access policies. |
| records | array[group_policy_object_central_access_policy] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "create_time": "2018-01-01 12:00:00 -0400",
      "description": "policy #1",
      "member_rules": [
        "r1",
        "r2"
      ],
      "name": "p1",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "update_time": "2018-01-01 12:00:00 -0400"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object_central_access_policy

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| description | string | Description about the policy. |
| member_rules | array[string] | Names of all central access rules applied to members. |
| name | string | |

| Name | Type | Description |
|---|---|---|
| sid | string | Security ID, unique identifier of the central policy. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an applied central access policy for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/central-access-policies/{name}`

**Introduced In:** 9.12

Retrieves applied central access policy for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy central-access-policy show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Central access policy name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| description | string | Description about the policy. |
| member_rules | array[string] | Names of all central access rules applied to members. |
| name | string | |
| sid | string | Security ID, unique identifier of the central policy. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

#### Example response

```
{
  "create_time": "2018-01-01 12:00:00 -0400",
  "description": "policy #1",
  "member_rules": [
    "r1",
    "r2"
  ],
  "name": "p1",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "update_time": "2018-01-01 12:00:00 -0400"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve applied central access rules for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/central-access-rules`

**Introduced In:** 9.12

Retrieves applied central access rules for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy central-access-rule show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| description | string | query | False | Filter by description |
| update_time | string | query | False | Filter by update_time |
| svm.name | string | query | False | Filter by svm.name |
| name | string | query | False | Filter by name<br>• minLength: 1 |
| proposed_permission | string | query | False | Filter by proposed_permission |
| resource_criteria | string | query | False | Filter by resource_criteria |
| create_time | string | query | False | Filter by create_time |
| current_permission | string | query | False | Filter by current_permission |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of central access rules. |
| records | array[group_policy_object_central_access_rule] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "create_time": "2018-01-01 12:00:00 -0400",
      "current_permission": "O:SYG:SYD:AR(A;;FA;;;WD)",
      "description": "rule #1",
      "name": "p1",
      "proposed_permission":
"O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)",
      "resource_criteria": "department",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "update_time": "2018-01-01 12:00:00 -0400"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object_central_access_rule

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| current_permission | string | Effective security policy in security descriptor definition language format. |
| description | string | Description about the policy. |
| name | string | |

| Name | Type | Description |
|---|---|---|
| proposed_permission | string | Proposed security policy in security descriptor definition language format. |
| resource_criteria | string | Criteria to scope resources for which access rules apply. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an applied central access rule for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/central-access-rules/{name}`

**Introduced In:** 9.12

Retrieves applied central access rule for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy central-access-rule show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Central access rule name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| create_time | string | Policy creation timestamp. |
| current_permission | string | Effective security policy in security descriptor definition language format. |
| description | string | Description about the policy. |
| name | string | |
| proposed_permission | string | Proposed security policy in security descriptor definition language format. |
| resource_criteria | string | Criteria to scope resources for which access rules apply. |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| update_time | string | Last policy modification timestamp. |

#### Example response


```
{
  "create_time": "2018-01-01 12:00:00 -0400",
  "current_permission": "O:SYG:SYD:AR(A;;FA;;;WD)",
  "description": "rule #1",
  "name": "p1",
  "proposed_permission":
"O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)",
  "resource_criteria": "department",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "update_time": "2018-01-01 12:00:00 -0400"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# View group policy objects

## View group policy objects

### Overview

When group policy objects (GPO) are enabled on your SMB server, ONTAP sends LDAP queries to the Active Directory server requesting GPO information.
If there are GPO definitions that are applicable to your SMB server, the Active Directory server returns the following GPO information:
• GPO name
• Current GPO version
• Location of the GPO definition
• Lists of UUIDs (universally unique identifiers) for GPO policy sets

## Retrieve applied group policy objects for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/objects`

**Introduced In:** 9.12

Retrieves applied group policy objects for specified SVM.

### Related ONTAP commands

- `vserver cifs group-policy show-applied`

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| registry_settings.branchcache.supported_hash_version | string | query | False | Filter by registry_settings.branchcache.supported_hash_version |
| registry_settings.branchcache.hash_publication_mode | string | query | False | Filter by registry_settings.branchcache.hash_publication_mode |
| registry_settings.refresh_time_random_offset | string | query | False | Filter by registry_settings.refresh_time_random_offset |
| registry_settings.refresh_time_interval | string | query | False | Filter by registry_settings.refresh_time_interval |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• minLength: 1 |
| svm.name | string | query | False | Filter by svm.name |
| central_access_policy_settings | string | query | False | Filter by central_access_policy_settings |
| file_system_path | string | query | False | Filter by file_system_path |
| enabled | boolean | query | False | Filter by enabled |
| central_access_policy_staging_audit_type | string | query | False | Filter by central_access_policy_staging_audit_type |
| ldap_path | string | query | False | Filter by ldap_path |
| extensions | string | query | False | Filter by extensions |
| security_settings.event_audit_settings.object_access_type | string | query | False | Filter by security_settings.event_audit_settings.object_access_type |
| security_settings.event_audit_settings.logon_type | string | query | False | Filter by security_settings.event_audit_settings.logon_type |
| security_settings.restrict_anonymous.no_enumeration_of_sam_accounts | boolean | query | False | Filter by security_settings.restrict_anonymous.no_enumeration_of_sam_accounts |
| security_settings.restrict_anonymous.combined_restriction_for_anonymous_user | string | query | False | Filter by security_settings.restrict_anonymous.combined_restriction_for_anonymous_user |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_settings.restrict_anonymous.no_enumeration_of_sam_accounts_and_shares | boolean | query | False | Filter by security_settings.restrict_anonymous.no_enumeration_of_sam_accounts_and_shares |
| security_settings.restrict_anonymous.anonymous_access_to_shares_and_named_pipes_restricted | boolean | query | False | Filter by security_settings.restrict_anonymous.anonymous_access_to_shares_and_named_pipes_restricted |
| security_settings.files_or_folders | string | query | False | Filter by security_settings.files_or_folders |
| security_settings.registry_values.signing_required | boolean | query | False | Filter by security_settings.registry_values.signing_required |
| security_settings.privilege_rights.change_notify_users | string | query | False | Filter by security_settings.privilege_rights.change_notify_users |
| security_settings.privilege_rights.security_privilege_users | string | query | False | Filter by security_settings.privilege_rights.security_privilege_users |
| security_settings.privilege_rights.take_ownership_users | string | query | False | Filter by security_settings.privilege_rights.take_ownership_users |
| security_settings.restricted_groups | string | query | False | Filter by security_settings.restricted_groups |
| security_settings.kerberos.max_ticket_age | string | query | False | Filter by security_settings.kerberos.max_ticket_age |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_settings.kerberos.max_renew_age | string | query | False | Filter by security_settings.kerberos.max_renew_age |
| security_settings.kerberos.max_clock_skew | string | query | False | Filter by security_settings.kerberos.max_clock_skew |
| security_settings.event_log_settings.max_size | integer | query | False | Filter by security_settings.event_log_settings.max_size |
| security_settings.event_log_settings.retention_method | string | query | False | Filter by security_settings.event_log_settings.retention_method |
| link | string | query | False | Filter by link |
| version | integer | query | False | Filter by version |
| index | integer | query | False | Filter by index |
| uuid | string | query | False | Filter by uuid |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of group policy objects. |
| records | array[group_policy_object] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "central_access_policy_settings": [
        "p1",
        "p2"
      ],
      "central_access_policy_staging_audit_type": "none",
      "extensions": [
        "audit",
        "security"
      ],
      "file_system_path": "\\test.com\\SysVol\\test.com\\policies
\\{42474212-3f9d-4489-ae01-6fcf4f805d4c}",
      "index": 1,
      "ldap_path": "cn={42474212-3f9d-4489-ae01-
6fcf4f805d4c},cn=policies,cn=system,DC=TEST,DC=COM",
      "link": "domain",
      "name": "test_policy",
      "registry_settings": {
        "branchcache": {
          "hash_publication_mode": "disabled",
          "supported_hash_version": "version1"
        },
        "refresh_time_interval": "P15M",
        "refresh_time_random_offset": "P1D"
      },
      "security_settings": {
        "event_audit_settings": {
          "logon_type": "failure",
          "object_access_type": "failure"
        },
        "event_log_settings": {
          "max_size": 2048,
          "retention_method": "do_not_overwrite"
```

```
        },
        "files_or_folders": [
          "/vol1/home",
          "/vol1/dir1"
        ],
        "kerberos": {
          "max_clock_skew": "P15M",
          "max_renew_age": "P2D",
          "max_ticket_age": "P24H"
        },
        "privilege_rights": {
          "change_notify_users": [
            "usr1",
            "usr2"
          ],
          "security_privilege_users": [
            "usr1",
            "usr2"
          ],
          "take_ownership_users": [
            "usr1",
            "usr2"
          ]
        },
        "restrict_anonymous": {
          "combined_restriction_for_anonymous_user": "no_access"
        },
        "restricted_groups": [
          "test_grp1",
          "test_grp2"
        ]
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "42474212-3f9d-4489-ae01-6fcf4f805d4c",
      "version": 7
    }
  ]
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

group_policy_object_branchcache

| Name | Type | Description |
|---|---|---|
| hash_publication_mode | string | Hash publication mode. |
| supported_hash_version | string | Hash version. |

group_policy_object_registry_setting

| Name | Type | Description |
|---|---|---|
| branchcache | group_policy_object_branchcache | |
| refresh_time_interval | string | Refresh time interval in ISO-8601 format. |
| refresh_time_random_offset | string | Random offset in ISO-8601 format. |

group_policy_object_event_audit

| Name | Type | Description |
|---|---|---|
| logon_type | string | Type of logon event to be audited. |
| object_access_type | string | Type of object access to be audited. |

group_policy_object_event_log

| Name | Type | Description |
|---|---|---|
| max_size | integer | Maximum size of security log, in kilobytes. |
| retention_method | string | Audit log retention method. |

group_policy_object_kerberos

| Name | Type | Description |
|---|---|---|
| max_clock_skew | string | Kerberos clock skew in ISO-8601 format. |
| max_renew_age | string | Kerberos max renew age in ISO-8601 format. |
| max_ticket_age | string | Kerberos max ticket age in ISO-8601 format. |

group_policy_object_privilege_right

| Name | Type | Description |
|---|---|---|
| change_notify_users | array[string] | Users with traversing bypass privileges. |
| security_privilege_users | array[string] | Users with security privileges. |
| take_ownership_users | array[string] | Users who can take ownership of securable objects. |

group_policy_object_registry_value

| Name | Type | Description |
|---|---|---|
| signing_required | boolean | SMB signing required. |

group_policy_object_restrict_anonymous

| Name | Type | Description |
|---|---|---|
| anonymous_access_to_shares_and_named_pipes_restricted | boolean | Restrict anonymous access to shares and named pipes. |
| combined_restriction_for_anonymous_user | string | Combined restriction for anonymous user. |

| Name | Type | Description |
|---|---|---|
| no_enumeration_of_sam_accounts | boolean | No enumeration of SAM accounts. |
| no_enumeration_of_sam_accounts_and_shares | boolean | No enumeration of SAM accounts and shares. |

group_policy_object_security_setting

| Name | Type | Description |
|---|---|---|
| event_audit_settings | group_policy_object_event_audit | |
| event_log_settings | group_policy_object_event_log | |
| files_or_folders | array[string] | Files/Directories for file security. |
| kerberos | group_policy_object_kerberos | |
| privilege_rights | group_policy_object_privilege_right | |
| registry_values | group_policy_object_registry_value | |
| restrict_anonymous | group_policy_object_restrict_anonymous | |
| restricted_groups | array[string] | List of restricted groups. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object

| Name | Type | Description |
| --- | --- | --- |
| central_access_policy_settings | array[string] | List of central access policies. |
| central_access_policy_staging_audit_type | string | Types of events to be audited. |
| enabled | boolean | Specifies whether group policies are enabled for the SVM. |
| extensions | array[string] | List of extensions. |
| file_system_path | string | File system path. |
| index | integer | Group policy object index. |
| ldap_path | string | LDAP path to the GPO. |
| link | string | Link info. |
| name | string | |
| registry_settings | group_policy_object_registry_setting | |
| security_settings | group_policy_object_security_setting | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| uuid | string | Policy UUID. |
| version | integer | Group policy object version. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an applied group policy object for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/objects/{index}`

**Introduced In:** 9.12

Retrieves applied group policy object for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | integer | path | True | Restricted group index. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| central_access_policy_settings | array[string] | List of central access policies. |
| central_access_policy_staging_audit_type | string | Types of events to be audited. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether group policies are enabled for the SVM. |
| extensions | array[string] | List of extensions. |
| file_system_path | string | File system path. |
| index | integer | Group policy object index. |
| ldap_path | string | LDAP path to the GPO. |
| link | string | Link info. |
| name | string | |
| registry_settings | group_policy_object_registry_setting | |
| security_settings | group_policy_object_security_setting | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| uuid | string | Policy UUID. |
| version | integer | Group policy object version. |

**Example response**


```
{
  "central_access_policy_settings": [
    "p1",
    "p2"
  ],
  "central_access_policy_staging_audit_type": "none",
  "extensions": [
    "audit",
    "security"
  ],
  "file_system_path": "\\test.com\\SysVol\\test.com\\policies
\\{42474212-3f9d-4489-ae01-6fcf4f805d4c}",
  "index": 1,
  "ldap_path": "cn={42474212-3f9d-4489-ae01-
6fcf4f805d4c},cn=policies,cn=system,DC=TEST,DC=COM",
  "link": "domain",
  "name": "test_policy",
  "registry_settings": {
    "branchcache": {
      "hash_publication_mode": "disabled",
      "supported_hash_version": "version1"
    },
    "refresh_time_interval": "P15M",
    "refresh_time_random_offset": "P1D"
  },
  "security_settings": {
    "event_audit_settings": {
      "logon_type": "failure",
      "object_access_type": "failure"
    },
    "event_log_settings": {
      "max_size": 2048,
      "retention_method": "do_not_overwrite"
    },
    "files_or_folders": [
      "/vol1/home",
      "/vol1/dir1"
    ],
    "kerberos": {
      "max_clock_skew": "P15M",
      "max_renew_age": "P2D",
      "max_ticket_age": "P24H"
    },
    "privilege_rights": {
```

```
      "change_notify_users": [
        "usr1",
        "usr2"
      ],
      "security_privilege_users": [
        "usr1",
        "usr2"
      ],
      "take_ownership_users": [
        "usr1",
        "usr2"
      ]
    },
    "restrict_anonymous": {
      "combined_restriction_for_anonymous_user": "no_access"
    },
    "restricted_groups": [
      "test_grp1",
      "test_grp2"
    ]
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "42474212-3f9d-4489-ae01-6fcf4f805d4c",
  "version": 7
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

group_policy_object_branchcache

| Name | Type | Description |
|---|---|---|
| hash_publication_mode | string | Hash publication mode. |
| supported_hash_version | string | Hash version. |

group_policy_object_registry_setting

| Name | Type | Description |
|---|---|---|
| branchcache | group_policy_object_branchcache | |
| refresh_time_interval | string | Refresh time interval in ISO-8601 format. |
| refresh_time_random_offset | string | Random offset in ISO-8601 format. |

group_policy_object_event_audit

| Name | Type | Description |
|---|---|---|
| logon_type | string | Type of logon event to be audited. |
| object_access_type | string | Type of object access to be audited. |

group_policy_object_event_log

| Name | Type | Description |
|---|---|---|
| max_size | integer | Maximum size of security log, in kilobytes. |
| retention_method | string | Audit log retention method. |

group_policy_object_kerberos

| Name | Type | Description |
|---|---|---|
| max_clock_skew | string | Kerberos clock skew in ISO-8601 format. |

| Name | Type | Description |
|---|---|---|
| max_renew_age | string | Kerberos max renew age in ISO-8601 format. |
| max_ticket_age | string | Kerberos max ticket age in ISO-8601 format. |

group_policy_object_privilege_right

| Name | Type | Description |
|---|---|---|
| change_notify_users | array[string] | Users with traversing bypass privileges. |
| security_privilege_users | array[string] | Users with security privileges. |
| take_ownership_users | array[string] | Users who can take ownership of securable objects. |

group_policy_object_registry_value

| Name | Type | Description |
|---|---|---|
| signing_required | boolean | SMB signing required. |

group_policy_object_restrict_anonymous

| Name | Type | Description |
|---|---|---|
| anonymous_access_to_shares_and_named_pipes_restricted | boolean | Restrict anonymous access to shares and named pipes. |
| combined_restriction_for_anonymous_user | string | Combined restriction for anonymous user. |
| no_enumeration_of_sam_accounts | boolean | No enumeration of SAM accounts. |
| no_enumeration_of_sam_accounts_and_shares | boolean | No enumeration of SAM accounts and shares. |

group_policy_object_security_setting

| Name | Type | Description |
|---|---|---|
| event_audit_settings | group_policy_object_event_audit | |
| event_log_settings | group_policy_object_event_log | |

| Name | Type | Description |
| --- | --- | --- |
| files_or_folders | array[string] | Files/Directories for file security. |
| kerberos | group_policy_object_kerberos | |
| privilege_rights | group_policy_object_privilege_right | |
| registry_values | group_policy_object_registry_value | |
| restrict_anonymous | group_policy_object_restrict_anonymous | |
| restricted_groups | array[string] | List of restricted groups. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve applied restricted group policies for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/restricted-groups`

**Introduced In:** 9.12

Retrieves applied policies of restricted groups for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy restricted-group show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| version | integer | query | False | Filter by version |
| link | string | query | False | Filter by link |
| policy_name | string | query | False | Filter by policy_name<br>• minLength: 1 |
| memberships | string | query | False | Filter by memberships |
| group_name | string | query | False | Filter by group_name<br>• minLength: 1 |
| members | string | query | False | Filter by members |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of restricted group info. |
| records | array[group_policy_object_restricted_group] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "group_name": "test_group",
      "link": "domain",
      "members": [
        "DOMAIN/test_user",
        "DOMAIN/user2"
      ],
      "memberships": [
        "DOMAIN/AdministratorGrp",
        "DOMAIN/deptMark"
      ],
      "policy_name": "test_policy",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "version": 7
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_policy_object_restricted_group

| Name | Type | Description |
|---|---|---|
| group_name | string | |
| link | string | Link info. |
| members | array[string] | Members of the group. |
| memberships | array[string] | Group is member of Group/OU. |
| policy_name | string | |

| Name | Type | Description |
|---|---|---|
| svm | svm | Will not be populated for objects that are yet to be applied. |
| version | integer | Group policy object version. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an applied restricted group policy for an SVM

GET `/protocols/cifs/group-policies/{svm.uuid}/restricted-groups/{policy_index}/{group_name}`

**Introduced In:** 9.12

Retrieves applied policy of restricted group for specified SVM.

#### Related ONTAP commands

- `vserver cifs group-policy restricted-group show-applied`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy_index | integer | path | True | Group policy index. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| group_name | string | path | True | Group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| group_name | string | |
| link | string | Link info. |
| members | array[string] | Members of the group. |
| memberships | array[string] | Group is member of Group/OU. |
| policy_name | string | |
| svm | svm | Will not be populated for objects that are yet to be applied. |
| version | integer | Group policy object version. |

#### Example response

```
{
  "group_name": "test_group",
  "link": "domain",
  "members": [
    "DOMAIN/test_user",
    "DOMAIN/user2"
  ],
  "memberships": [
    "DOMAIN/AdministratorGrp",
    "DOMAIN/deptMark"
  ],
  "policy_name": "test_policy",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "version": 7
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Will not be populated for objects that are yet to be applied.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage home directory search paths

## Manage home directory search paths

### Overview

ONTAP home directory functionality can be used to create home directories for SMB users on the CIFS server and automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user.

The home directory search path is a set of absolute paths from the root of an SVM that directs ONTAP to search for home directories. If there are multiple search paths, ONTAP tries them in the order specified until it finds a valid path. To use the CIFS home directories feature, at least one home directory search path must be added for an SVM.

### Examples

## Creating a home directory search path

To create a home directory search path, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/home-directory/search-paths

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-
paths?return_records=true" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"path\": \"/\", \"svm\": { \"name\": \"vs1\",
\"uuid\": \"a41fd873-ecf8-11e8-899d-0050568e9333\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "a41fd873-ecf8-11e8-899d-0050568e9333",
      "name": "vs1"
    },
  "path": "/"
  }
]
}
```

**Retrieving the CIFS home directory search paths configuration for all SVMs in the cluster**

```
# The API:
GET /protocols/cifs/home-directory/search-paths

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-
paths?fields=*&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
      "name": "vs1"
    },
    "index": 1,
    "path": "/"
  },
  {
    "svm": {
      "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
      "name": "vs1"
    },
      "index": 2,
      "path": "/a"
  },
  {
    "svm": {
      "uuid": "4f23449b-f4ce-11e8-b075-0050568e278e",
      "name": "vs2"
    },
    "index": 1,
    "path": "/"
  },
  {
    "svm": {
      "uuid": "4f23449b-f4ce-11e8-b075-0050568e278e",
      "name": "vs2"
    },
    "index": 2,
    "path": "/1"
  }
],
```

```
"num_records": 4
}
```

#### Retrieving a specific home directory searchpath configuration for an SVM

The configuration returned is identified by the UUID of its SVM and the index (position) in the list of search paths that is searched to find a home directory of a user.

```
# The API:
GET /api/protocols/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-
paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
  "name": "vs1"
},
"index": 2,
"path": "/a"
}
```

#### Reordering a specific home directory search path in the list

An entry in the home directory search path list can be reordered to a new positin by specifying the 'new_index' field. The reordered configuration is identified by the UUID of its SVM and the index.

```
# The API:
PATCH /api/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-
paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2?new_index=1" -H "accept:
application/json"
```

#### Removing a specific home directory search path for an SVM

The entry being removed is identified by the UUID of its SVM and the index.

```
# The API:
DELETE /api/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/home-
directory/search-paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2" -H "accept:
application/json"
```

### Retrieve CIFS home directory search paths

GET `/protocols/cifs/home-directory/search-paths`

**Introduced In:** 9.6

Retrieves CIFS home directory search paths.

#### Related ONTAP commands

- `cifs server home-directory search-path show`

#### Learn more

- DOC /protocols/cifs/home-directory/search-paths

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | query | False | Filter by path |
| index | integer | query | False | Filter by index |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_search_path] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "index": 0,
      "path": "/HomeDirectory/EngDomain",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_search_path

This is a list of CIFS home directory search paths. When a CIFS client connects to a home directory share, these paths are searched in the order indicated by the position field to find the home directory of the connected CIFS client.

| Name | Type | Description |
|---|---|---|
| index | integer | The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST. |

| Name | Type | Description |
|---|---|---|
| path | string | The file system path that is searched to find the home directory of the CIFS client. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a home directory search path

POST `/protocols/cifs/home-directory/search-paths`

**Introduced In:** 9.6

Creates a home directory search path.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the home directory search path.
- `path` - Path in the owning SVM namespace that is used to search for home directories.

#### Related ONTAP commands

- `cifs server home-directory search-path add`

#### Learn more

- DOC /protocols/cifs/home-directory/search-paths

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| index | integer | The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST. |
| path | string | The file system path that is searched to find the home directory of the CIFS client. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "index": 0,
  "path": "/HomeDirectory/EngDomain",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[cifs_search_path] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "index": 0,
      "path": "/HomeDirectory/EngDomain",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655551 | Invalid home-directory search-path path |
| 655462 | The specified path is an invalid file-type |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_search_path

This is a list of CIFS home directory search paths. When a CIFS client connects to a home directory share, these paths are searched in the order indicated by the position field to find the home directory of the connected CIFS client.

| Name | Type | Description |
|---|---|---|
| index | integer | The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST. |
| path | string | The file system path that is searched to find the home directory of the CIFS client. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a CIFS home directory search path

DELETE `/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}`

**Introduced In:** 9.6

Deletes a CIFS home directory search path.

#### Related ONTAP commands

- `cifs server home-directory search-path remove`

#### Learn more

- DOC /protocols/cifs/home-directory/search-paths

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | integer | path | True | Home directory search path index |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a CIFS home directory search path for an SVM

GET `/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}`

**Introduced In:** 9.6

Retrieves a CIFS home directory search path of an SVM.

#### Related ONTAP commands

- `cifs server home-directory search-path show`

#### Learn more

- DOC /protocols/cifs/home-directory/search-paths

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | integer | path | True | Home directory search path index |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| index | integer | The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST. |
| path | string | The file system path that is searched to find the home directory of the CIFS client. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "index": 0,
  "path": "/HomeDirectory/EngDomain",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Re-order a CIFS home directory search path

PATCH `/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}`

**Introduced In:** 9.6

Reorders a CIFS home directory search path.

#### Related ONTAP commands

- `cifs server home-directory search-path reorder`

#### Learn more

- DOC /protocols/cifs/home-directory/search-paths

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | integer | path | True | Home directory search path index |
| new_index | integer | query | False | New position for the home directory search path |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655463 | Failed to reorder the search-path because the new-index is invalid. It cannot be '0' and it cannot go beyond the current entries |

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage local groups

### Manage local groups

#### Overview

The CIFS server can use local groups for authorization when determining share, file and directory access rights.
You can use this API to display local group information and to control local group configurations.

#### Retrieving local group information

The local group GET endpoint retrieves all of the local groups configurations for data SVMs.

#### Examples

### Retrieving all of the fields for all of the local group configurations

The local group GET endpoint retrieves all of the local groups configurations for data SVMs.

```
# The API:
/api/protocols/cifs/local-groups

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-groups?fields=**"
-H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-32-544",
    "name": "BUILTIN\\Administrators",
    "description": "Built-in Administrators group",
    "members": [
      {
        "name": "CIFS_SERVER\\Administrator"
      },
      {
        "name": "AD_DOMAIN\\Domain Admins"
      }
    ]
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-32-545",
    "name": "BUILTIN\\Users",
    "description": "All users",
    "members": [
      {
        "name": "AD_DOMAIN\\Domain Users"
      }
    ]
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-32-546",
    "name": "BUILTIN\\Guests",
    "description": "Built-in Guests Group",
    "members": [
```

```
      {
        "name": "SACHILDAP02\\Domain Guests"
      }
    ]
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-32-547",
    "name": "BUILTIN\\Power Users",
    "description": "Restricted administrative privileges"
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-32-551",
    "name": "BUILTIN\\Backup Operators",
    "description": "Backup Operators group"
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
    "name": "CIFS_SERVER\\group2",
    "description": "local group2"
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
      "name": "vs2"
    },
    "sid": "S-1-5-32-544",
    "name": "BUILTIN\\Administrators",
    "description": "Built-in Administrators group",
    "members": [
      {
        "name": "VS2.CIFS\\Administrator"
      },
      {
        "name": "VS2.CIFS\\user3"
```

```
      },
      {
        "name": "SACHILDAP02\\Domain Admins"
      }
    ]
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
      "name": "vs2"
    },
    "sid": "S-1-5-32-545",
    "name": "BUILTIN\\Users",
    "description": "All users",
    "members": [
      {
        "name": "SACHILDAP02\\Domain Users"
      }
    ]
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
      "name": "vs2"
    },
    "sid": "S-1-5-32-546",
    "name": "BUILTIN\\Guests",
    "description": "Built-in Guests Group",
    "members": [
      {
        "name": "SACHILDAP02\\Domain Guests"
      }
    ]
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
      "name": "vs2"
    },
    "sid": "S-1-5-32-547",
    "name": "BUILTIN\\Power Users",
    "description": "Restricted administrative privileges"
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
```

```
      "name": "vs2"
    },
    "sid": "S-1-5-32-551",
    "name": "BUILTIN\\Backup Operators",
    "description": "Backup Operators group"
  },
  {
    "svm": {
      "uuid": "5060077c-5be6-11eb-90b7-0050568e5169",
      "name": "vs2"
    },
    "sid": "S-1-5-21-1625922807-3304708894-3529444428-1001",
    "name": "CIFS_SERVER\\group1",
    "description": "local group1"
  }
],
"num_records": 12
}
```


**Retrieving a local group configuration of a specific SVM and group**

```
# The API:
/api/protocols/cifs/local-groups/{svm.uuid}/{sid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-groups/25b363a6-
2971-11eb-88e1-0050568eefd4/S-1-5-21-256008430-3394229847-3930036330-1001"
-H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
  "name": "vs1"
},
"sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
"name": "CIFS_SERVER\\group1",
"description": "local group"
}
```

#### Creating a local group configuration

The local group POST endpoint creates a local group configuration for the specified SVM.

#### Example

```
# The API:
/api/protocols/cifs/local-groups

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/local-groups" -H
"accept: application/json" -H "Content-Type: application/json" -d '{
"svm": {"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1"}, "name":
"group1"}'
```

#### Updating a local group configuration

The local group PATCH endpoint updates the name and description of the specified local group and the specified SVM.

#### Example

### Update the local group name from 'group1' to 'group2'

```
# The API:
/api/protocols/local-groups/{svm.uuid}/{sid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/local-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/S-1-5-21-256008430-3394229847-3930036330-1257"
-H "accept: application/json" -H "Content-Type: application/json" -d '{
"name": "group2", "description": "local group"}'
```

#### Deleting a local group configuration

The local group DELETE endpoint deletes the specified local group of the specified SVM.

#### Example

### Delete the local group 'group1'

```
# The API:
/api/protocols/cifs/local-groups/{svm.uuid}/{sid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/local-
groups/179d3c85-7053-11e8-b9b8-005056b41bd1/S-1-5-21-256008430-3394229847-
3930036330-1001" -H "accept: application/json"
```

### Retrieve local groups for all SVMs

GET `/protocols/cifs/local-groups`

**Introduced In:** 9.9

Retrieves the local groups for all of the SVMs.

#### Advanced properties

- `members`

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group show`
- `vserver cifs users-and-groups local-group show-members`

#### Learn more

- DOC /protocols/cifs/local-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | query | False | Filter by sid<br>• Introduced in: 9.10 |
| members.name | string | query | False | Filter by members.name |
| name | string | query | False | Filter by name<br>• minLength: 1 |
| description | string | query | False | Filter by description<br>• maxLength: 256 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of local group records. |
| records | array[local_cifs_group] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "This is a local group",
      "members": [
        {
          "name": "string"
        }
      ],
      "name": "SMB_SERVER01\\group",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

members

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| description | string | Description for the local group. |
| members | array[members] | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create the local group configuration for an SVM**

POST `/protocols/cifs/local-groups`

**Introduced In:** 9.9

Creates the local group configuration for the specified SVM.

#### Important notes

- The group name can contain up to 256 characters.
- The group name cannot be terminated by a period.
- The group name does not support any of the following characters: " / ? [ ] , : \| < > + = ; ? * @ or ASCII characters in the range 1-31.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the local group.
- `name` - Name of the local group.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group create`

#### Learn more

- DOC /protocols/cifs/local-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| description | string | Description for the local group. |
| members | array[members] | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "description": "This is a local group",
  "members": [
    {
      "name": "string"
    }
  ],
  "name": "SMB_SERVER01\\group",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262278 | Name is a required field. |
| 655399 | CIFS server must exist to create a local group. |
| 655660 | The operation is allowed only on data SVMs. |
| 655661 | The group name and description should not exceed 256 characters. |
| 655668 | The specified group name contains illegal characters. |

| Error Code | Description |
| --- | --- |
| 655675 | The local domain name specified in the group name does not exist. |
| 655677 | This operation does not allow for the creation of a group in the BUILTIN domain. |
| 655682 | The group name cannot be blank. |
| 655717 | The specified group name already exists. |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

members

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group

| Name | Type | Description |
|---|---|---|
| description | string | Description for the local group. |
| members | array[members] | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a local group configuration for an SVM

DELETE `/protocols/cifs/local-groups/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Deletes a local group configuration for the specified SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group delete`

#### Learn more

- DOC /protocols/cifs/local-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | path | True | Local group SID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve local group information for a group and SVM

GET `/protocols/cifs/local-groups/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Retrieves local group information for the specified group and SVM.

#### Advanced properties

- `members`

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group show`
- `vserver cifs users-and-groups local-group show-members`

#### Learn more

- DOC /protocols/cifs/local-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | path | True | Local group SID |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| description | string | Description for the local group. |
| members | array[members] | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "This is a local group",
  "members": [
    {
      "name": "string"
    }
  ],
  "name": "SMB_SERVER01\\group",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

members

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update local group information for a group in an SVM

PATCH `/protocols/cifs/local-groups/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Updates the local group information of the specified group in the specified SVM. This API can also be used to rename a local group.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group modify`
- `vserver cifs users-and-groups local-group rename`

#### Learn more

- DOC /protocols/cifs/local-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | path | True | Local group SID |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| description | string | Description for the local group. |
| members | array[members] | |

| Name | Type | Description |
|---|---|---|
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "description": "This is a local group",
  "members": [
    {}
  ],
  "name": "SMB_SERVER01\\group",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655661 | The group name and description should not exceed 256 characters. |

| Error Code | Description |
| --- | --- |
| 655668 | The specified group name contains illegal characters. |
| 655675 | The local domain name specified in the group name does not exist. |
| 655682 | The group name cannot be blank. |
| 655712 | To rename an existing group, the local domain specified in name must match the local domain of the group to be renamed. |
| 655713 | Failed to rename a group. The error code returned details the failure along with the reason for the failure. Take corrective actions as per the specified reason. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

members

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group

| Name | Type | Description |
|---|---|---|
| description | string | Description for the local group. |
| members | array[members] | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage local group members

### Manage local group members

#### Overview

You can use this API to display local group members and to add or delete local users, Active Directory users and/or Active Directory groups to a local group of an SVM.

#### Examples

### Retrieving the members of a specific local group

```
# The API:
/api/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-groups/2ebec9c7-
28be-11eb-95f4-0050568ed0a2/S-1-5-21-256008430-3394229847-3930036330-
1257/members" -H "accept: application/json"

# The response:
{
"records": [
  {
    "name": "CIFS_SERVER1\\user1"
  },
  {
    "name": "CIFS_SERVER1\\user2"
  }
],
"num_records": 2
}
```

#### Adding members to a local group

The local group members POST endpoint adds local users, Active Directory users and/or Active Directory groups to the specified local group and the SVM.

### Adding local users to a group

```
# The API:
/api/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/local-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/S-1-5-21-256008430-3394229847-3930036330-
1001/members" -H "accept: application/json" -H "Content-Type:
application/json" -d '{ "records": [ { "name": "user1" }, { "name":
"user2"} ] }'
```

#### Deleting local users from the local group of a specific SVM

#### Example

### Delete the local users 'user1' and 'user2' from the specified local group

```
# The API:
/api/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/local-
groups/179d3c85-7053-11e8-b9b8-005056b41bd1/S-1-5-21-256008430-3394229847-
3930036330-1001/members" -H "accept: application/json" -d '{ "records": [
{ "name": "user1"}, { "name": "user2"} ] }'
```

### Delete local users or Active Directory users and groups from a local group and SVM

DELETE `/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members`

**Introduced In:** 9.10

Deletes the local users, Active Directory users and/or Active Directory groups from the specified local group and SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group remove-members`

#### Learn more

- DOC /protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| local_cifs_group.sid | string | path | True | Local group SID |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |
| records | array[records] | An array of local users, Active Directory users, and Active Directory groups specified to add or delete multiple members to or from a local group in a single API call. Not allowed when the `name` property is used. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "local_cifs_group": {
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
  },
  "name": "string",
  "records": [
    {
      "name": "string"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error ontap error response codes | error code |
description | | ---------- | ----------- | | 655673     | failed to
resolve the member to be deleted from the specified group. | | 655719
| failed to delete a member from the specified group. the error code
returned details the failure along with the reason for the failure. take
corrective actions as per the specified reason. | | 655742     | records
must not be specified when member name is specified. | | 655743     | svm
uuid and cifs local group sid are invalid fields for the "records"
parameter. |
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

local_cifs_group

| Name | Type | Description |
|---|---|---|
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

records

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group_members

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |
| records | array[records] | An array of local users, Active Directory users, and Active Directory groups specified to add or delete multiple members to or from a local group in a single API call.<br>Not allowed when the `name` property is used. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve local users and Active Directory users and groups that belong to a local group and SVM**

GET `/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members`

**Introduced In:** 9.10

Retrieves local users, Active Directory users and Active Directory groups which are members of the specified

local group and SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group show-members`

#### Learn more

- DOC /protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| local_cifs_group.sid | string | path | True | Local group SID |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | svm.uuid | string | path | True |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| UUID of the SVM to which this object belongs. | fields | array[string] | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of local group records. |
| records | array[local_cifs_group_members] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "local_cifs_group": {
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
      },
      "name": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

local_cifs_group

| Name | Type | Description |
|---|---|---|
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

records

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group_members

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add local users and Active Directory users and groups to a local group and SVM

POST `/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members`

**Introduced In:** 9.10

Adds local users, Active Directory users and Active Directory groups to the specified local group and SVM.

#### Important note

- Specified members are appended to the existing list of members.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM for which members are added to local group.
- `local_cifs_group.sid` - Security ID of the local group to which members are added.
- `name` or `records` - Local users, Active Directory users, or Active Directory groups to be added to a particular local group.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group add-members`

#### Learn more

- DOC /protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| local_cifs_group.sid | string | path | True | Local group SID |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

| Name | Type | Description |
|---|---|---|
| records | array[records] | An array of local users, Active Directory users, and Active Directory groups specified to add or delete multiple members to or from a local group in a single API call. Not allowed when the `name` property is used. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "local_cifs_group": {
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
  },
  "name": "string",
  "records": [
    {
      "name": "string"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655673 | Failed to resolve the member to be added to the specified group. |
| 655718 | Failed to add a member to the specified group. The error code returned details the failure along with the reason for the failure. Take corrective actions as per the specified reason. |
| 655742 | Records must not be specified when member name is specified. |
| 655743 | SVM UUID and CIFS local group SID are invalid fields for the "records" parameter. |
| 655744 | Either name or records must be specified. |

**Definitions**

**See Definitions**

local_cifs_group

| Name | Type | Description |
|---|---|---|
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

records

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_group_members

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |
| records | array[records] | An array of local users, Active Directory users, and Active Directory groups specified to add or delete multiple members to or from a local group in a single API call.<br>Not allowed when the `name` property is used. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete a local user, or Active Directory user or group from a local group and SVM**

DELETE `/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members/{name}`

**Introduced In:** 9.10

Deletes the local user, Active Directory user and/or Active Directory group from the specified local group and

SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group remove-members`

#### Learn more

- DOC /protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| local_cifs_group.sid | string | path | True | Local group SID |
| name | string | path | True | Member name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error ontap error response codes | error code |
description | | ---------- | ----------- | | 655673     | failed to
resolve the member to be deleted from the specified group. | | 655719
| failed to delete a member from the specified group. the error code
returned details the failure along with the reason for the failure. take
corrective actions as per the specified reason. |
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a local user, Active Directory user or group belonging to a local group and SVM

GET `/protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members/{name}`

**Introduced In:** 9.10

Retrieves local user, Active Directory user and Active Directory group which is member of the specified local group and SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-group show-members`

#### Learn more

- DOC /protocols/cifs/local-groups/{svm.uuid}/{local_cifs_group.sid}/members

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| local_cifs_group.sid | string | path | True | Local group SID |
| name | string | path | True | Member name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| local_cifs_group | local_cifs_group | |
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response

```
{
  "local_cifs_group": {
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
  },
  "name": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

local_cifs_group

| Name | Type | Description |
|---|---|---|
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

records

| Name | Type | Description |
|---|---|---|
| name | string | Local user, Active Directory user, or Active Directory group which is a member of the specified local group. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage local users

### Manage local users

#### Overview

The CIFS server can use local users for CIFS authentication. The local users can also be used for authorization when determining both share and file/directory access rights to data residing on the storage virtual machine (SVM).
You can use this API to display local user information and to control local user configurations.

#### Retrieving local user information

You can use the local user GET endpoint to retrieve all of the local user configurations for data SVMs.

#### Examples

### Retrieving all of the fields for local user configurations for all SVMs

```
# The API:
/api/protocols/cifs/local-users

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-users?fields=**"
-H "accept: application/json"

# The response:
```

```
{
"records": [
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-500",
    "name": "CIFS_SERVER1\\Administrator",
    "full_name": "",
    "description": "Built-in administrator account",
    "account_disabled": false,
    "membership": [
      {
        "name": "BUILTIN\\Administrators",
        "sid": "S-1-5-32-544"
      }
    ]
  },
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
    "name": "CIFS_SERVER1\\user1",
    "full_name": "local user1",
    "description": "This is CIFS local user",
    "account_disabled": false
  },
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1002",
    "name": "CIFS_SERVER1\\user2",
    "full_name": "local user2",
    "description": "This is CIFS local user",
    "account_disabled": false,
    "membership": [
      {
        "name": "CIFS_SERVER1\\grp1",
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
      },
      {
```

```
        "name": "CIFS_SERVER1\\grp2",
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1002"
      }
    ]
  },
  {
    "svm": {
      "uuid": "3f479a01-2971-11eb-88e1-0050568eefd4",
      "name": "vs2"
    },
    "sid": "S-1-5-21-1625922807-3304708894-3529444428-500",
    "name": "CIFS_SERVER2\\Administrator",
    "full_name": "",
    "description": "Built-in administrator account",
    "account_disabled": false,
    "membership": [
      {
        "name": "BUILTIN\\Administrators",
        "sid": "S-1-5-32-544"
      }
    ]
  },
  {
    "svm": {
      "uuid": "3f479a01-2971-11eb-88e1-0050568eefd4",
      "name": "vs2"
    },
    "sid": "S-1-5-21-1625922807-3304708894-3529444428-1001",
    "name": "CIFS_SERVER2\\user1",
    "full_name": "local user1",
    "description": "This is CIFS local user",
    "account_disabled": false
  }
],
"num_records": 5
}
```


#### Retrieving the local user configuration of a specific SVM

```
# The API:
/api/protocols/cifs/local-users
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-
users?svm.uuid=25b363a6-2971-11eb-88e1-0050568eefd4&fields=**" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-500",
    "name": "CIFS_SERVER1\\Administrator",
    "full_name": "",
    "description": "Built-in administrator account",
    "account_disabled": false,
    "membership": [
      {
        "name": "BUILTIN\\Administrators",
        "sid": "S-1-5-32-544"
      }
    ]
  },
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
    "name": "CIFS_SERVER1\\user1",
    "full_name": "local user1",
    "description": "This is CIFS local user",
    "account_disabled": false
  },
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "sid": "S-1-5-21-256008430-3394229847-3930036330-1002",
    "name": "CIFS_SERVER1\\user2",
    "full_name": "local user2",
    "description": "This is CIFS local user",
    "account_disabled": false,
```

```
    "membership": [
      {
        "name": "CIFS_SERVER1\\grp1",
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
      },
      {
        "name": "CIFS_SERVER1\\grp2",
        "sid": "S-1-5-21-256008430-3394229847-3930036330-1002"
      }
    ]
  },
],
"num_records": 3
}
```

#### Retrieving a local user configuration of a specific SVM and user

```
# The API:
/api/protocols/cifs/local-users/{svm.uuid}/{sid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/local-users/25b363a6-
2971-11eb-88e1-0050568eefd4/S-1-5-21-1625922807-3304708894-3529444428-
1001" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
  "name": "vs1"
},
"sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
"name": "CIFS_SERVER1\\user1",
"full_name": "local user1",
"description": "This is CIFS local user",
"account_disabled": false
}
```

#### Creating a local user configuration

The local user POST endpoint creates a local user configuration for the specified SVM.

#### Examples

### Creating a local user configuration with all fields specified

```
# The API:
/api/protocols/cifs/local-users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/local-users" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"account_disabled\": false, \"description\": \"This is local user.\",
\"full_name\": \"user name\", \"name\": \"SMB_SERVER01\\\\user\",
\"password\": \"netapp1N\", \"svm\": { \"name\": \"vs1\", \"uuid\":
\"25b363a6-2971-11eb-88e1-0050568eefd4\" }}"
```

### Creating a local user configuration with only mandatory fields specified

```
# The API:
/api/protocols/cifs/local-users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/local-users" -H
"accept: application/json" -d "{ \"name\": \"user1\", \"password\":
\"netapp1N\", \"svm\": { \"uuid\": \"25b363a6-2971-11eb-88e1-
0050568eefd4\" }}"
```

#### Updating a local user configuration

The local user PATCH endpoint updates the local user configuration for the specified user and SVM.

### Updating a local user name and password

```
# The API:
/api/protocols/cifs/local-users/{svm.uuid}/{sid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/local-users/25b363a6-
2971-11eb-88e1-0050568eefd4/S-1-5-21-1625922807-3304708894-3529444428-
1001" -H "accept: application/json" -d "{ \"name\": \"new_user1\",
\"password\": \"netapp1Net\" }"
```

---

#### Deleting a local user configuration

The local user DELETE endpoint deletes the specified local user for the specified SVM. The following example shows a DELETE operation:

```
# The API:
/api/protocols/cifs/local-users/{svm.uuid}/{sid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/local-users/25b363a6-
2971-11eb-88e1-0050568eefd4/S-1-5-21-1625922807-3304708894-3529444428-
1001" -H "accept: application/json"
```

---

### Retrieve local users for all SVMs

GET `/protocols/cifs/local-users`

**Introduced In:** 9.9

Retrieves local users for all of the SVMs. Local groups to which this user belongs to are also displayed.

#### Advanced properties

- `membership`

#### Related ONTAP commands

- `vserver cifs users-and-groups local-user show`
- `vserver cifs users-and-groups local-user show-membership`

#### Learn more

- DOC /protocols/cifs/local-users

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• minLength: 1 |
| description | string | query | False | Filter by description<br>• maxLength: 256 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| sid | string | query | False | Filter by sid<br>• Introduced in: 9.10 |
| full_name | string | query | False | Filter by full_name<br>• maxLength: 256 |
| account_disabled | boolean | query | False | Filter by account_disabled |
| membership.sid | string | query | False | Filter by membership.sid<br>• Introduced in: 9.10 |
| membership.name | string | query | False | Filter by membership.name<br>• minLength: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of local user records. |
| records | array[local_cifs_user] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "This is local user.",
      "full_name": "user name",
      "membership": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "SMB_SERVER01\\group",
          "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
        }
      ],
      "name": "SMB_SERVER01\\user",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

membership

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_user

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |
| membership | array[membership] | Specifies local groups of which this local user is a member. |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the local user configuration for an SVM

POST `/protocols/cifs/local-users`

**Introduced In:** 9.9

Creates the local user configuration for an SVM.

#### Important notes

- The user name can contain up to 20 characters.
- The user name cannot be terminated by a period.
- The user name does not support any of the following characters: " / ? [ ] , : \| < > + = ; ? * @ or ASCII characters in the range 1-31.
- The password must be at least six characters in length and must not contain the user account name.
- The password must contain characters from three of the following four categories:
- English uppercase characters (A through Z)
- English lowercase characters (a through z)
- Base 10 digits (0 through 9)
- Special characters: ~ ! @ # 0 ^ & * _ - + = ` ? \| ( ) [ ] : ; " ' < > , . ? /

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the local user.
- `name` - Name of the local user.
- `password` - Password for the local user.

#### Default property values

If not specified in POST, the following default property value is assigned:

- `account_disabled` - false

#### Related ONTAP commands

- `vserver cifs users-and-groups local-user create`

#### Learn more

- DOC /protocols/cifs/local-users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |
| membership | array[membership] | Specifies local groups of which this local user is a member. |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| password | string | Password for the local user. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "description": "This is local user.",
  "full_name": "user name",
  "membership": [
    {
      "name": "SMB_SERVER01\\group",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
    }
  ],
  "name": "SMB_SERVER01\\user",
  "password": "string",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of local user records. |
| records | array[local_cifs_user] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "description": "This is local user.",
      "full_name": "user name",
      "membership": [
        {
          "name": "SMB_SERVER01\\group",
          "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
        }
      ],
      "name": "SMB_SERVER01\\user",
      "password": "string",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262278 | Name and password are required fields. |
| 655399 | CIFS server must exist to create a local user. |
| 655660 | The operation is allowed only on data SVMs. |

| Error Code | Description |
| --- | --- |
| 655661 | The user name should not exceed 20 characters. Also full_name and description should not exceed 256 characters. |
| 655665 | The user name must not match the CIFS server name of the specified SVM. |
| 655668 | The specified user name contains illegal characters. |
| 655675 | The local domain name specified in user name doesn’t exist. |
| 655682 | The user name cannot be blank. |
| 655733 | The password does not meet the password complexity requirements. |
| 655736 | The specified user name already exists. |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

membership

| Name | Type | Description |
|---|---|---|
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_user

| Name | Type | Description |
|---|---|---|
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |

| Name | Type | Description |
|---|---|---|
| membership | array[membership] | Specifies local groups of which this local user is a member. |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| password | string | Password for the local user. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete the local user configuration for an SVM**

DELETE `/protocols/cifs/local-users/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Deletes a local user configuration for the specified SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-user delete`

#### Learn more

- DOC /protocols/cifs/local-users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | path | True | Local user SID |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655735 | The local Administrator account cannot be deleted. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the local user information for a user and SVM

GET `/protocols/cifs/local-users/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Retrieves local user information for the specified user and SVM.

#### Advanced properties

- `membership`

#### Related ONTAP commands

- `vserver cifs users-and-groups local-user show`
- `vserver cifs users-and-groups local-user show-membership`

#### Learn more

- DOC /protocols/cifs/local-users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| sid | string | path | True | Local user SID |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |
| membership | array[membership] | Specifies local groups of which this local user is a member. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "This is local user.",
  "full_name": "user name",
  "membership": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "SMB_SERVER01\\group",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
    }
  ],
  "name": "SMB_SERVER01\\user",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

membership

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the local user information for a user and SVM

PATCH `/protocols/cifs/local-users/{svm.uuid}/{sid}`

**Introduced In:** 9.10

Updates local user information for the specified user and SVM. The PATCH endpoint is also used to rename a user and to set the password for the user.

#### Related ONTAP commands

- `vserver cifs users-and-groups local-user modify`
- `vserver cifs users-and-groups local-user rename`
- `vserver cifs users-and-groups local-user set-password`

#### Learn more

- DOC /protocols/cifs/local-users

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| sid | string | path | True | Local user SID |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |
| membership | array[membership] | Specifies local groups of which this local user is a member. |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| password | string | Password for the local user. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "description": "This is local user.",
  "full_name": "user name",
  "membership": [
    {
      "name": "SMB_SERVER01\\group",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001"
    }
  ],
  "name": "SMB_SERVER01\\user",
  "password": "string",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655661 | The user name should not exceed 20 characters. Also full_name and description should not exceed 256 characters. |
| 655668 | The specified user name contains illegal characters. |
| 655675 | The local domain name specified in the user name does not exist. |
| 655682 | The user name cannot be blank. |
| 655732 | Failed to rename a user. The error code returned details the failure along with the reason for the failure. Take corrective actions as per the specified reason. |

| Error Code | Description |
| --- | --- |
| 655733 | The specified password does not meet the password complexity requirements. |
| 655737 | To rename an existing user, the local domain specified in name must match the local domain of the user to be renamed. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

membership

| Name | Type | Description |
|---|---|---|
| name | string | Local group name. The maximum supported length of a group name is 256 characters. |
| sid | string | The security ID of the local group which uniquely identifies the group. The group SID is automatically generated in POST and it is retrieved using the GET method. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_user

| Name | Type | Description |
|---|---|---|
| account_disabled | boolean | Indicates whether the user account is enabled or disabled. |
| description | string | Description for the local user. |
| full_name | string | User's full name. |

| Name | Type | Description |
|---|---|---|
| membership | array[membership] | Specifies local groups of which this local user is a member. |
| name | string | Local user name. The maximum supported length of an user name is 20 characters. |
| password | string | Password for the local user. |
| sid | string | The security ID of the local user which uniquely identifies the user. The user SID is automatically generated in POST and it is retrieved using the GET method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View NetBIOS information

### View NetBIOS information

Displays information about NBT connections for the cluster. Displays the IP address associated with the interfaces, the IP addresses of the WINS servers in use, and information about the registered NetBIOS names

for the cluster. You can use this command to troubleshoot NetBIOS name resolution problems.

#### Examples

### Retrieving full CIFS NetBIOS information for all SVMs in the cluster

```
# The API:
GET /api/protocols/cifs/netbios

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/netbios?return_timeout=15&fields=*&return_records=t
rue" -H  "accept: application/json" -H  "authorization: Basic
<CREDENTIALS>"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "c2179c2c-777f-11ec-95a1-1315568ef5fd",
      "name": "vsim2"
    },
    "svm": {
      "name": "vs1",
      "uuid": "99a3bd71-777f-11ec-95a1-1315568ef5fd"
    },
    "name": "CIFSERVER2",
    "interfaces": [
      "172.10.144.44"
    ],
    "wins_servers": [
      {
        "ip": "10.10.10.10",
        "state": "inactive"
      }
    ],
    "scope": "group",
    "mode": "h",
    "state": "broadcast",
    "time_left": 0,
    "name_registration_type": "",
    "suffix": "00"
  },
  {
```

```
    "node": {
      "uuid": "c2179c2c-777f-11ec-95a1-1315568ef5fd",
      "name": "vsim2"
    },
    "svm": {
      "name": "vs1",
      "uuid": "99a3bd71-777f-11ec-95a1-1315568ef5fd"
    },
    "name": "CIFSERVER2",
    "interfaces": [
      "172.10.144.44"
    ],
    "wins_servers": [
      {
        "ip": "10.10.10.10",
        "state": "inactive"
      }
    ],
    "scope": "",
    "mode": "h",
    "state": "broadcast",
    "time_left": 0,
    "name_registration_type": "group",
    "suffix": "20"
  }
],
"num_records": 2
}
```


### Retrieving CIFS NetBIOS Information for a particular SVM

```
# The API:
GET /api/protocols/cifs/netbios

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/netbios?svm.uuid=45a3bd71-777f-11ec-95a1-
1315568ef5fd&return_timeout=15&fields=*&return_records=true" -H  "accept:
application/json" -H  "authorization: Basic <CREDENTIALS>"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "c2179c2c-777f-11ec-95a1-1315568ef5fd",
      "name": "vsim2"
    },
    "svm": {
      "name": "vs1",
      "uuid": "45a3bd71-777f-11ec-95a1-1315568ef5fd"
    },
    "name": "CIFSERVER2",
    "interfaces": [
      "172.10.144.44"
    ],
    "wins_servers": [
      {
        "ip": "10.10.10.10",
        "state": "inactive"
      }
    ],
    "scope": "group",
    "mode": "h",
    "state": "broadcast",
    "time_left": 0,
    "name_registration_type": "",
    "suffix": "00"
  }
],
"num_records": 1
}
```

### Retrieve NetBIOS information

GET `/protocols/cifs/netbios`

**Introduced In:** 9.11

Retrieves NetBIOS information.

#### Related ONTAP commands

- `vserver cifs nbtstat`

#### Learn more

- DOC /protocols/cifs/netbios

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| mode | string | query | False | Filter by mode |
| name_registration_type | string | query | False | Filter by name_registration_type |
| time_left | integer | query | False | Filter by time_left |
| scope | string | query | False | Filter by scope |
| wins_servers.state | string | query | False | Filter by wins_servers.state |
| wins_servers.ip | string | query | False | Filter by wins_servers.ip |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name |
| state | string | query | False | Filter by state |
| suffix | string | query | False | Filter by suffix |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interfaces | string | query | False | Filter by interfaces |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of Records. |
| records | array[netbios] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "interfaces": [
        [
          "10.10.10.35"
        ]
      ],
      "mode": "string",
      "name": "CLUSTER_1",
      "name_registration_type": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": "string",
      "state": "string",
      "suffix": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "time_left": 0,
      "wins_servers": [
        {
```

```
            "ip": "10.224.65.20",
            "state": "string"
          }
        ]
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

server

| Name | Type | Description |
| --- | --- | --- |
| ip | string | Windows Internet Name Server (WINS) address which manages and maps the NetBIOS name of the CIFS server to their network IP addresses. The IP addresses are IPv4 addresses. |
| state | string | Specifies the state of the WINS server. |

netbios

| Name | Type | Description |
| --- | --- | --- |
| interfaces | array[string] | |
| mode | string | Specifies the mode in which the NetBIOS name service is configured. The supported values are:<br>• p - Point to Point<br>• h - Hybrid<br>• m - Mixed<br>• b - Broadcast |
| name | string | Specifies the NetBIOS name. |
| name_registration_type | string | Specifies the name registration type. |
| node | node | |
| scope | string | Specifies the NetBIOS name scope. Scope is used as a name for the set of NetBIOS nodes that participate in a NetBIOS over TCP (NBT) virtual LAN. |
| state | string | Specifies the registration state of the NetBIOS Name. |

| Name | Type | Description |
| --- | --- | --- |
| suffix | string | Specifies the NetBIOS suffix. A NetBIOS suffix is the 16th character of the 16-character NetBIOS name. The NetBIOS suffix is used by Microsoft Networking software to identify functionality installed on the registered device. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| time_left | integer | Specifies the registration time left with WINS, in minutes. |
| wins_servers | array[server] | List of WINS |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage CIFS services

### Manage CIFS services

#### Overview

A CIFS server is necessary to provide SMB clients with access to the Storage Virtual Machine (SVM). Before you begin, the following prerequisites must be in place:

- At least one SVM LIF must exist on the SVM.
- The LIFs must be able to connect to the DNS servers configured on the SVM and to an Active Directory domain controller of the domain to which you want to join the CIFS server.
- The DNS servers must contain the service location records that are needed to locate the Active Directory domain services.
- The cluster time must be synchronized to within five minutes of the Active Directory domain controller.

#### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

### Information on the CIFS server

You must keep the following in mind when creating the CIFS server:

- The CIFS server name might or might not be the same as the SVM name.
- The CIFS server name can be up to 15 characters in length.
- The following characters are not allowed: @ # * ( ) = + [ ] \| ; : " , < > \ / ?
- You must use the FQDN when specifying the domain.
- The default is to add the CIFS server machine account to the Active Directory "CN=Computer" object.
- You can choose to add the CIFS server to a different organizational unit (OU) by specifying the "organizational_unit" parameter. When specifying the OU, do not specify the domain portion of the distinguished name; only specify the OU or CN portion of the distinguished name. ONTAP appends the value provided for the required "-domain" parameter onto the value provided for the "-ou" parameter to create the Active Directory distinguished name, which is used when joining the Active Directory domain.
- You can optionally choose to add a text comment of up to 256 characters about the CIFS server. If there is a space in the comment text, you must enclose the entire string in quotation marks.
- You can optionally choose to add a comma-delimited list of one or more NetBIOS aliases for the CIFS server.
- The initial administrative status of the CIFS server is "up".
- The large-mtu feature is enabled for the new CIFS server.
- The multichannel feature is disabled for the new CIFS server.
- If LDAP is configured with the *use_start_tls* and *session_security* features, the new CIFS server will also have this property set.
- The security.kdc_encryption and security.advertised_kdc_encryptions fields are mutually exclusive. Use the advertised_kdc_encryptions field to specify the encryption types to be advertised to the Key Distribution Center (KDC) server in the Active Directory domain.

#### Examples

### Creating a CIFS server

To create a CIFS server, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/services

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/services?return_timeout=10&return_records=true" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{ \"ad_domain\": { \"fqdn\":
\"ontapavc.com\", \"organizational_unit\": \"CN=Computers\", \"password\":
\"<PASSWORD-FOR-USER>\", \"user\": \"administrator\" }, \"comment\":
\"This CIFS Server Belongs to CS Department\", \"default_unix_user\":
\"string\", \"enabled\": true, \"name\": \"CIFS1\", \"netbios\": {
\"aliases\": [ \"ALIAS_1\", \"ALIAS_2\", \"ALIAS_3\" ], \"enabled\":
false, \"wins_servers\": [ \"10.224.65.20\", \"10.224.65.21\" ] },
\"options\": { \"admin_to_root_mapping\": true, \"advanced_sparse_file\":
true, \"copy_offload\": true, \"fake_open\": true, \"fsctl_trim\": true,
\"junction_reparse\": true, \"large_mtu\": true, \"multichannel\": true,
\"null_user_windows_name\": \"string\", \"path_component_cache\": true,
\"referral\": false, \"smb_credits\": 128, \"widelink_reparse_versions\":
[ \"smb1\" ] }, \"security\": { \"encrypt_dc_connection\": false,
\"kdc_encryption\": false, \"restrict_anonymous\": \"no_enumeration\",
\"session_security\": \"none\", \"smb_encryption\": false,
\"smb_signing\": false, \"use_ldaps\": false, \"use_start_tls\": false },
\"svm\": { \"name\": \"vs1\", \"uuid\": \"e0c20d9c-96cd-11eb-97da-
0050568e684d\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "e0c20d9c-96cd-11eb-97da-0050568e684d",
      "name": "vs1"
    },
    "name": "CIFS1",
    "ad_domain": {
      "fqdn": "ONTAPAVC.COM",
      "organizational_unit": "CN=Computers"
    },
    "enabled": true,
    "comment": "This CIFS Server Belongs to CS Department",
    "security": {
      "restrict_anonymous": "no_enumeration",
      "smb_signing": false,
```

```
      "smb_encryption": false,
      "kdc_encryption": false,
      "aes_netlogon_enabled": false,
      "try_ldap_channel_binding": false,
      "referral_enabled": false,
      "lm_compatibility_level": "lm_ntlm_ntlmv2_krb",
      "encrypt_dc_connection": false,
      "use_start_tls": false,
      "session_security": "none",
      "use_ldaps": false
    },
    "netbios": {
      "aliases": [
        "ALIAS_1",
        "ALIAS_2",
        "ALIAS_3"
      ],
      "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
      ],
      "enabled": false
    },
    "default_unix_user": "string",
    "options": {
      "advanced_sparse_file": true,
      "referral": false,
      "widelink_reparse_versions": [
        "smb1"
      ],
      "multichannel": true,
      "path_component_cache": true,
      "null_user_windows_name": "string",
      "junction_reparse": true,
      "fsctl_trim": true,
      "large_mtu": true,
      "fake_open": true,
      "smb_credits": 128,
      "admin_to_root_mapping": true,
      "copy_offload": true
    }
  }
],
"job": {
  "uuid": "825a0b4b-9703-11eb-8cc1-0050568e684d",
  "_links": {
```

```
    "self": {
      "href": "/api/cluster/jobs/825a0b4b-9703-11eb-8cc1-0050568e684d"
    }
  }
}
}
```

#### Creating a CIFS server when using AKV

```
# The API:
POST /api/protocols/cifs/services

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/services?return_timeout=10&return_records=true" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{ \"key_vault_uri\":
\"https://testkv.vault.azure.net\", \"client_secret\": \"<CLIENT-SECRET-
ACCESS-KEY>\", \"authentication_method\": \"client_secret\",
\"tenant_id\": \"c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe\", \"client_id\":
\"e959d1b5-5a63-4284-9268-851e30e3eceb\", \"ad_domain\": { \"fqdn\":
\"ontapavc.com\", \"organizational_unit\": \"CN=Computers\", \"user\":
\"administrator\" }, \"comment\": \"This CIFS Server Belongs to CS
Department\", \"default_unix_user\": \"string\", \"enabled\": true,
\"name\": \"CIFS1\", \"netbios\": { \"aliases\": [ \"ALIAS_1\",
\"ALIAS_2\", \"ALIAS_3\" ], \"enabled\": false, \"wins_servers\": [
\"10.224.65.20\", \"10.224.65.21\" ] }, \"options\": {
\"admin_to_root_mapping\": true, \"advanced_sparse_file\": true,
\"copy_offload\": true, \"fake_open\": true, \"fsctl_trim\": true,
\"junction_reparse\": true, \"large_mtu\": true, \"multichannel\": true,
\"null_user_windows_name\": \"string\", \"path_component_cache\": true,
\"referral\": false, \"smb_credits\": 128, \"widelink_reparse_versions\":
[ \"smb1\" ] }, \"security\": { \"encrypt_dc_connection\": false,
\"kdc_encryption\": false, \"restrict_anonymous\": \"no_enumeration\",
\"session_security\": \"none\", \"smb_encryption\": false,
\"smb_signing\": false, \"use_ldaps\": false, \"use_start_tls\": false },
\"svm\": { \"name\": \"vs1\", \"uuid\": \"e0c20d9c-96cd-11eb-97da-
0050568e684d\" }}"

# The response:
{
"num_records": 1,
```

```
"records": [
  {
    "svm": {
      "uuid": "e0c20d9c-96cd-11eb-97da-0050568e684d",
      "name": "vs1"
    },
    "name": "CIFS1",
    "ad_domain": {
      "fqdn": "ONTAPAVC.COM",
      "organizational_unit": "CN=Computers"
    },
    "enabled": true,
    "comment": "This CIFS Server Belongs to CS Department",
    "security": {
      "restrict_anonymous": "no_enumeration",
      "smb_signing": false,
      "smb_encryption": false,
      "kdc_encryption": false,
      "aes_netlogon_enabled": false,
      "try_ldap_channel_binding": false,
      "referral_enabled": false,
      "lm_compatibility_level": "lm_ntlm_ntlmv2_krb",
      "encrypt_dc_connection": false,
      "use_start_tls": false,
      "session_security": "none",
      "use_ldaps": false
    },
    "netbios": {
      "aliases": [
        "ALIAS_1",
        "ALIAS_2",
        "ALIAS_3"
      ],
      "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
      ],
      "enabled": false
    },
    "default_unix_user": "string",
    "options": {
      "advanced_sparse_file": true,
      "referral": false,
      "widelink_reparse_versions": [
        "smb1"
      ],
```

```
      "multichannel": true,
      "path_component_cache": true,
      "null_user_windows_name": "string",
      "junction_reparse": true,
      "fsctl_trim": true,
      "large_mtu": true,
      "fake_open": true,
      "smb_credits": 128,
      "admin_to_root_mapping": true,
      "copy_offload": true
    }
  }
],
"job": {
  "uuid": "825a0b4b-9703-11eb-8cc1-0050568e684d",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/825a0b4b-9703-11eb-8cc1-0050568e684d"
    }
  }
}
}
```


### Creating a CIFS server for Hybrid user

```
# The API:
POST /api/protocols/cifs/services

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/services?return_timeout=10&return_records=true" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{ \"auth_user_type\":
\"hybrid_user\", \"client_certificate\": \"certificate added to azure
application\", \"authentication_method\": \"certificate\", \"tenant_id\":
\"c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe\", \"client_id\": \"e959d1b5-5a63-
4284-9268-851e30e3eceb\", \"ad_domain\": { \"fqdn\": \"ontapavc.com\",
\"organizational_unit\": \"CN=Computers\" }, \"comment\": \"This CIFS
Server Belongs to CS Department\", \"enabled\": true, \"name\": \"CIFS1\",
\"svm\": { \"name\": \"vs1\", \"uuid\": \"e0c20d9c-96cd-11eb-97da-
0050568e684d\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "e0c20d9c-96cd-11eb-97da-0050568e684d",
      "name": "vs1"
    },
    "name": "CIFS1",
    "ad_domain": {
      "fqdn": "ONTAPAVC.COM",
      "organizational_unit": "CN=Computers"
    },
    "enabled": true,
    "comment": "This CIFS Server Belongs to CS Department",
  }
],
"job": {
  "uuid": "825a0b4b-9703-11eb-8cc1-0050568e684d",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/825a0b4b-9703-11eb-8cc1-0050568e684d"
    }
  }
}
}
```

**Creating a CIFS server when using the security.advertised_kdc_encryptions field**

```
# The API:
POST /api/protocols/cifs/services

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/services?return_timeout=120&return_records=true&for
ce=true" -H "accept: application/json" -H "authorization: Basic
<CREDENTIALS>" -H "Content-Type: application/json" -d "{ \"ad_domain\": {
\"fqdn\": \"testdomain.com\", \"organizational_unit\": \"CN=Computers\",
\"password\": \"<PASSWORD-FOR-USER>\", \"user\": \"administrator\" },
\"enabled\": true, \"name\": \"CIFS_server1\", \"security\": {
\"advertised_kdc_encryptions\": [ \"aes-128\" ] }, \"svm\": { \"name\":
\"testvs\", \"uuid\": \"508375d4-fb16-11ee-a792-005056a75b5c\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "508375d4-fb16-11ee-a792-005056a75b5c",
      "name": "testvs"
    },
    "name": "CIFS_SERVER1",
    "domain_workgroup": "TESTDOMAIN",
    "ad_domain": {
      "fqdn": "TESTDOMAIN.COM",
      "organizational_unit": "CN=Computers",
      "default_site": ""
    },
    "enabled": true,
    "comment": "",
    "security": {
      "restrict_anonymous": "no_enumeration",
      "smb_signing": false,
      "smb_encryption": false,
      "kdc_encryption": true,
      "lm_compatibility_level": "lm_ntlm_ntlmv2_krb",
      "aes_netlogon_enabled": false,
      "try_ldap_channel_binding": true,
      "ldap_referral_enabled": false,
      "advertised_kdc_encryptions": [
```

```
        "aes_128"
      ],
      "encrypt_dc_connection": false,
      "use_start_tls": false,
      "session_security": "none",
      "use_ldaps": false
    },
    "netbios": {
      "wins_servers": [],
      "enabled": false
    },
    "default_unix_user": "pcuser",
    "statistics": {
      "timestamp": "2024-04-16T15:17:41Z",
      "status": "ok",
      "throughput_raw": {
        "read": 0,
        "write": 0,
        "total": 0
      },
      "iops_raw": {
        "read": 0,
        "write": 0,
        "other": 0,
        "total": 0
      },
      "latency_raw": {
        "read": 0,
        "write": 0,
        "other": 0,
        "total": 0
      }
    },
    "options": {
      "advanced_sparse_file": true,
      "referral": false,
      "widelink_reparse_versions": [
        "smb1"
      ],
      "multichannel": false,
      "path_component_cache": true,
      "junction_reparse": true,
      "fsctl_trim": true,
      "large_mtu": true,
      "fake_open": true,
      "smb_credits": 512,
```

```
      "admin_to_root_mapping": true,
      "copy_offload": true,
      "shadowcopy": true,
      "shadowcopy_dir_depth": 5,
      "export_policy_enabled": false
    },
    "group_policy_object_enabled": false
  }
],
"job": {
  "uuid": "73eec230-fc04-11ee-a792-005056a75b5c",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/73eec230-fc04-11ee-a792-005056a75b5c"
    }
  }
}
}
```


#### Retrieving the full CIFS server configuration for all SVMs in the cluster

```
# The API:
GET /api/protocols/cifs/services

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/services?fields=*&return_records=true&return_timeou
t=15" -H "accept: application/json" -H "authorization: Basic
<CREDENTIALS>"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "e0c20d9c-96cd-11eb-97da-0050568e684d",
      "name": "vs1"
    },
    "name": "CIFS1",
    "ad_domain": {
      "fqdn": "ONTAPAVC.COM",
      "organizational_unit": "CN=Computers"
```

```
    },
    "enabled": true,
    "comment": "This CIFS Server Belongs to CS Department",
    "security": {
      "restrict_anonymous": "no_enumeration",
      "smb_signing": false,
      "smb_encryption": false,
      "kdc_encryption": false,
      "aes_netlogon_enabled": false,
      "try_ldap_channel_binding": false,
      "referral_enabled": false,
      "lm_compatibility_level": "lm_ntlm_ntlmv2_krb",
      "encrypt_dc_connection": false,
      "use_start_tls": false,
      "session_security": "none",
      "use_ldaps": false
    },
    "netbios": {
      "aliases": [
        "ALIAS_1",
        "ALIAS_2",
        "ALIAS_3"
      ],
      "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
      ],
      "enabled": false
    },
    "default_unix_user": "string",
    "options": {
      "advanced_sparse_file": true,
      "referral": false,
      "widelink_reparse_versions": [
        "smb1"
      ],
      "multichannel": true,
      "path_component_cache": true,
      "null_user_windows_name": "string",
      "junction_reparse": true,
      "fsctl_trim": true,
      "large_mtu": true,
      "fake_open": true,
      "smb_credits": 128,
      "admin_to_root_mapping": true,
      "copy_offload": true
```

```
    }
  }
],
"num_records": 1
}
```

#### Retrieving CIFS server configuration details for a specific SVM

```
# The API:
GET /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/services/e0c20d9c-96cd-
11eb-97da-0050568e684d" -H "accept: application/json" -H "authorization:
Basic <CREDENTIALS>"

# The response:
{
"svm": {
  "uuid": "e0c20d9c-96cd-11eb-97da-0050568e684d",
  "name": "vs1"
},
"name": "CIFS1",
"ad_domain": {
  "fqdn": "ONTAPAVC.COM",
  "organizational_unit": "CN=Computers"
},
"enabled": true,
"comment": "This CIFS Server Belongs to CS Department",
"security": {
  "restrict_anonymous": "no_enumeration",
  "smb_signing": false,
  "smb_encryption": false,
  "kdc_encryption": false,
  "aes_netlogon_enabled": false,
  "try_ldap_channel_binding": false,
  "referral_enabled": false,
  "lm_compatibility_level": "lm_ntlm_ntlmv2_krb",
  "encrypt_dc_connection": false,
  "use_start_tls": false,
  "session_security": "none",
  "use_ldaps": false
```

```
},
"netbios": {
  "aliases": [
    "ALIAS_1",
    "ALIAS_2",
    "ALIAS_3"
  ],
  "wins_servers": [
    "10.224.65.20",
    "10.224.65.21"
  ],
  "enabled": false
},
"default_unix_user": "string",
"options": {
  "advanced_sparse_file": true,
  "referral": false,
  "widelink_reparse_versions": [
    "smb1"
  ],
  "multichannel": true,
  "path_component_cache": true,
  "null_user_windows_name": "string",
  "junction_reparse": true,
  "fsctl_trim": true,
  "large_mtu": true,
  "fake_open": true,
  "smb_credits": 128,
  "admin_to_root_mapping": true,
  "copy_offload": true
}
}
```


### Updating CIFS server properties for the specified SVM

```
# The API:
PATCH /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/services/e0c20d9c-
96cd-11eb-97da-0050568e684d" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"comment\": \"CIFS SERVER MODIFICATION\"}"
```

---

#### Removing a CIFS server for a specific SVM

To delete a CIFS server, use the following API. This will delete the CIFS server along with other CIFS configurations such as CIFS share, share ACLs, homedir search-path, and so on.

---

```
# The API:
DELETE /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/services/e0c20d9c-
96cd-11eb-97da-0050568e684d" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"ad_domain\": { \"fqdn\": \"ontapavc.com\",
\"organizational_unit\": \"CN=Computers\", \"password\": \"<PASSWORD-FOR-
USER>\", \"user\": \"administrator\" }, \"force\": true}"

# The call when using AKV:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/services/e0c20d9c-
96cd-11eb-97da-0050568e684d" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"key_vault_uri\": \"https://testkv.vault.azure.net\",
\"client_secret\": \"<AWS-SECRET-ACCESS-KEY>\", \"authentication_method\":
\"client_secret\", \"tenant_id\": \"c9f32fcb-4ab7-40fe-af1b-
1850d46cfbbe\", \"client_id\": \"e959d1b5-5a63-4284-9268-851e30e3eceb\",
\"ad_domain\": { \"fqdn\": \"ontapavc.com\", \"organizational_unit\":
\"CN=Computers\" , \"user\": \"administrator\" }, \"force\": true}"

# The call when using EntraId:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/services/e0c20d9c-
96cd-11eb-97da-0050568e684d" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"auth_user_type\": \"hybrid-user\", \"client_certificate\":
\"client_certificate_added_to_azure_application\",
\"authentication_method\": \"certificate\", \"tenant_id\": \"c9f32fcb-
4ab7-40fe-af1b-1850d46cfbbe\", \"client_id\": \"e959d1b5-5a63-4284-9268-
851e30e3eceb\", \"ad_domain\": { \"fqdn\": \"ontapavc.com\",
\"organizational_unit\": \"CN=Computers\"}, \"force\": true}"
```

---

### Retrieve CIFS servers

GET `/protocols/cifs/services`

**Introduced In:** 9.6

Retrieves CIFS servers.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific

fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver cifs server show`
- `vserver cifs server options show`
- `vserver cifs server security show`

#### Learn more

- DOC /protocols/cifs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| proxy_username | string | query | False | Filter by proxy_username<br>• Introduced in: 9.15 |
| auth_user_type | string | query | False | Filter by auth_user_type<br>• Introduced in: 9.16 |
| security.restrict_anonymous | string | query | False | Filter by security.restrict_anonymous |
| security.use_ldaps | boolean | query | False | Filter by security.use_ldaps<br>• Introduced in: 9.10 |
| security.smb_signing | boolean | query | False | Filter by security.smb_signing |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security.ldap_referral_enabled | boolean | query | False | Filter by security.ldap_referral_enabled<br>• Introduced in: 9.10 |
| security.kdc_encryption | boolean | query | False | Filter by security.kdc_encryption |
| security.aes_netlogon_enabled | boolean | query | False | Filter by security.aes_netlogon_enabled<br>• Introduced in: 9.10 |
| security.session_security | string | query | False | Filter by security.session_security<br>• Introduced in: 9.10 |
| security.try_ldap_channel_binding | boolean | query | False | Filter by security.try_ldap_channel_binding<br>• Introduced in: 9.10 |
| security.use_start_tls | boolean | query | False | Filter by security.use_start_tls<br>• Introduced in: 9.10 |
| security.lm_compatibility_level | string | query | False | Filter by security.lm_compatibility_level<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security.encrypt_dc_connection | boolean | query | False | Filter by security.encrypt_dc_connection<br>• Introduced in: 9.8 |
| security.smb_encryption | boolean | query | False | Filter by security.smb_encryption |
| security.advertised_kdc_encryptions | string | query | False | Filter by security.advertised_kdc_encryptions<br>• Introduced in: 9.12 |
| proxy_host | string | query | False | Filter by proxy_host<br>• Introduced in: 9.15 |
| client_id | string | query | False | Filter by client_id<br>• Introduced in: 9.15 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| ad_domain.default_site | string | query | False | Filter by ad_domain.default_site<br>• Introduced in: 9.13 |
| ad_domain.fqdn | string | query | False | Filter by ad_domain.fqdn |
| ad_domain.organizational_unit | string | query | False | Filter by ad_domain.organizational_unit |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| workgroup | string | query | False | Filter by workgroup<br>• Introduced in: 9.15<br>• maxLength: 15<br>• minLength: 1 |
| tenant_id | string | query | False | Filter by tenant_id<br>• Introduced in: 9.15 |
| verify_host | boolean | query | False | Filter by verify_host<br>• Introduced in: 9.15 |
| oauth_host | string | query | False | Filter by oauth_host<br>• Introduced in: 9.15 |
| options.widelink_reparse_versions | string | query | False | Filter by options.widelink_reparse_versions<br>• Introduced in: 9.10 |
| options.shadowcopy | boolean | query | False | Filter by options.shadowcopy<br>• Introduced in: 9.11 |
| options.max_same_user_sessions_per_connection | integer | query | False | Filter by options.max_same_user_sessions_per_connection<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| options.max_watches_set_per_tree | integer | query | False | Filter by options.max_watches_set_per_tree<br>• Introduced in: 9.16 |
| options.max_opens_same_file_per_tree | integer | query | False | Filter by options.max_opens_same_file_per_tree<br>• Introduced in: 9.16 |
| options.copy_offload | boolean | query | False | Filter by options.copy_offload<br>• Introduced in: 9.10 |
| options.fake_open | boolean | query | False | Filter by options.fake_open<br>• Introduced in: 9.10 |
| options.referral | boolean | query | False | Filter by options.referral<br>• Introduced in: 9.10 |
| options.max_same_tree_connect_per_session | integer | query | False | Filter by options.max_same_tree_connect_per_session<br>• Introduced in: 9.16 |
| options.smb_credits | integer | query | False | Filter by options.smb_credits<br>• Introduced in: 9.10<br>• Max value: 8192<br>• Min value: 2 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| options.advanced_sparse_file | boolean | query | False | Filter by options.advanced_sparse_file<br><br>• Introduced in: 9.10 |
| options.dac_enabled | boolean | query | False | Filter by options.dac_enabled<br><br>• Introduced in: 9.15 |
| options.junction_reparse | boolean | query | False | Filter by options.junction_reparse<br><br>• Introduced in: 9.10 |
| options.shadowcopy_dir_depth | integer | query | False | Filter by options.shadowcopy_dir_depth<br><br>• Introduced in: 9.11 |
| options.multichannel | boolean | query | False | Filter by options.multichannel<br><br>• Introduced in: 9.10 |
| options.fsctl_trim | boolean | query | False | Filter by options.fsctl_trim<br><br>• Introduced in: 9.10 |
| options.trusted_domain_enum_search_enabled | boolean | query | False | Filter by options.trusted_domain_enum_search_enabled<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| options.max_lifs_per_session | integer | query | False | Filter by options.max_lifs_per_session<br>• Introduced in: 9.16<br>• Max value: 256<br>• Min value: 1 |
| options.client_dup_detection_enabled | boolean | query | False | Filter by options.client_dup_detection_enabled<br>• Introduced in: 9.16 |
| options.path_component_cache | boolean | query | False | Filter by options.path_component_cache<br>• Introduced in: 9.10 |
| options.admin_to_root_mapping | boolean | query | False | Filter by options.admin_to_root_mapping<br>• Introduced in: 9.10 |
| options.client_version_reporting_enabled | boolean | query | False | Filter by options.client_version_reporting_enabled<br>• Introduced in: 9.16 |
| options.null_user_windows_name | string | query | False | Filter by options.null_user_windows_name<br>• Introduced in: 9.10 |
| options.large_mtu | boolean | query | False | Filter by options.large_mtu<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| options.max_connections_per_session | integer | query | False | Filter by options.max_connections_per_session<br>• Introduced in: 9.16<br>• Max value: 1024<br>• Min value: 2 |
| options.backup_symlink_enabled | boolean | query | False | Filter by options.backup_symlink_enabled<br>• Introduced in: 9.15 |
| options.export_policy_enabled | boolean | query | False | Filter by options.export_policy_enabled<br>• Introduced in: 9.13 |
| proxy_type | string | query | False | Filter by proxy_type<br>• Introduced in: 9.15 |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| default_unix_user | string | query | False | Filter by default_unix_user |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br>• Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.7 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.7 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.7 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.7 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.7 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.7 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.7 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.7 |
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.7 |
| authentication_method | string | query | False | Filter by authentication_method<br><br>• Introduced in: 9.15 |
| timeout | integer | query | False | Filter by timeout<br><br>• Introduced in: 9.15 |
| group_policy_object_enabled | boolean | query | False | Filter by group_policy_object_enabled<br><br>• Introduced in: 9.12 |
| netbios.enabled | boolean | query | False | Filter by netbios.enabled |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| netbios.wins_servers | string | query | False | Filter by netbios.wins_servers |
| netbios.aliases | string | query | False | Filter by netbios.aliases<br>• maxLength: 15<br>• minLength: 1 |
| auth-style | string | query | False | Filter by auth-style<br>• Introduced in: 9.15 |
| name | string | query | False | Filter by name<br>• maxLength: 15<br>• minLength: 1 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br>• Introduced in: 9.7 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br>• Introduced in: 9.7 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br>• Introduced in: 9.7 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.7 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.7 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.7 |
| proxy_port | integer | query | False | Filter by proxy_port<br><br>• Introduced in: 9.15 |
| enabled | boolean | query | False | Filter by enabled |
| key_vault_uri | string | query | False | Filter by key_vault_uri<br><br>• Introduced in: 9.15 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_domain": {
        "default_site": "string",
        "fqdn": "example.com",
        "organizational_unit": "string"
      },
      "auth-style": "domain",
      "auth_user_type": "string",
      "authentication_method": "string",
      "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
      "comment": "This CIFS Server Belongs to CS Department",
      "default_unix_user": "string",
      "key_vault_uri": "https://kmip-akv-keyvault.vault.azure.net/",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
```

```
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "CIFS1",
      "netbios": {
        "aliases": [
          "ALIAS_1",
          "ALIAS_2",
          "ALIAS_3"
        ],
        "wins_servers": [
          "10.224.65.20",
          "10.224.65.21"
        ]
      },
      "oauth_host": "login.microsoftonline.com",
      "options": {
        "max_connections_per_session": 32,
        "max_lifs_per_session": 256,
        "max_opens_same_file_per_tree": 1000,
        "max_same_tree_connect_per_session": 5000,
        "max_same_user_sessions_per_connection": 2500,
        "max_watches_set_per_tree": 500,
        "null_user_windows_name": "string",
        "smb_credits": 128,
        "widelink_reparse_versions": [
          "smb1"
        ]
      },
      "proxy_host": "proxy.eng.com",
      "proxy_port": 1234,
      "proxy_type": "string",
      "proxy_username": "proxyuser",
      "security": {
        "advertised_kdc_encryptions": [
          "string"
        ],
        "lm_compatibility_level": "string",
        "restrict_anonymous": "string",
```

```
        "session_security": "string"
      },
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
      "timeout": 25,
      "workgroup": "workgrp1"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

cifs_netbios

| Name | Type | Description |
| --- | --- | --- |
| aliases | array[string] | |
| enabled | boolean | Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration. |
| wins_servers | array[string] | |

cifs_service_options

| Name | Type | Description |
| --- | --- | --- |
| admin_to_root_mapping | boolean | Specifies whether or not Administrator can be mapped to the UNIX user "root". |
| advanced_sparse_file | boolean | Specifies whether or not the CIFS server supports the advanced sparse file capabilities. This allows<br>CIFS clients to query the allocated ranges of a file and to write zeroes or free data blocks for ranges<br>of a file. |
| backup_symlink_enabled | boolean | Specifies whether or not to preserve UNIX symlinks during backup through SMB. |
| client_dup_detection_enabled | boolean | Specifies whether or not client duplicate session detection is enabled for CIFS. |

| Name | Type | Description |
| --- | --- | --- |
| client_version_reporting_enabled | boolean | Specifies whether or not client version reporting is enabled for CIFS. |
| copy_offload | boolean | Specifies whether or not to enable the Copy Offload feature. This feature enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.<br><br>Note that this will also enable/disable the direct copy feature accordingly. |
| dac_enabled | boolean | Specifies whether or not the Dynamic Access Control (DAC) feature is enabled for the CIFS server. |
| export_policy_enabled | boolean | Specifies whether or not export policies are enabled for CIFS. |
| fake_open | boolean | Specifies whether or not fake open support is enabled. This parameter allows you to optimize the<br>open and close requests coming from SMB 2 clients. |
| fsctl_trim | boolean | Specifies whether or not the trim requests (FSCTL_FILE_LEVEL_TRIM) are supported on the CIFS server. |
| junction_reparse | boolean | Specifies whether or not the reparse point support is enabled. When enabled the CIFS server exposes junction points to Windows clients as reparse points. This parameter is only active<br>if the client has negotiated use of the SMB 2 or SMB 3 protocol. This parameter is not supported for SVMs with Infinite Volume. |

| Name | Type | Description |
| --- | --- | --- |
| large_mtu | boolean | Specifies whether or not SMB clients can send reads up to 1 MB in size. |
| max_connections_per_session | integer | Specifies the maximum number of connections allowed per multichannel session. |
| max_lifs_per_session | integer | Specifies the maximum number of LIFs advertised per multichannel session. |
| max_opens_same_file_per_tree | integer | Specifies the maximum number of opens on the same file per tree. |
| max_same_tree_connect_per_session | integer | Specifies the maximum number of same tree connections per session. |
| max_same_user_sessions_per_connection | integer | Specifies the maximum number of same user sessions per connection. |
| max_watches_set_per_tree | integer | Specifies the maximum number of watches set per tree. |
| multichannel | boolean | Specifies whether or not the CIFS server supports Multichannel. |
| null_user_windows_name | string | Specifies a Windows User or Group name that should be mapped in case of a NULL user value. |
| path_component_cache | boolean | Specifies whether or not the path component cache is enabled on the CIFS server. |
| referral | boolean | Specifies whether or not to refer clients to more optimal LIFs. When enabled, it automatically refers clients to a data LIF local to the node which hosts the root of the requested share. |

| Name | Type | Description |
| --- | --- | --- |
| shadowcopy | boolean | Specifies whether or not to enable the Shadowcopy Feature. This feature enables to take share-based backup copies of data that is in a data-consistent state at a specific point in time where the data is accessed over SMB 3.0 shares. |
| shadowcopy_dir_depth | integer | Specifies the maximum level of subdirectories on which ONTAP should create shadow copies.<br>• Default value: 5<br>• Introduced in: 9.11<br>• x-nullable: true |
| smb_credits | integer | Specifies the maximum number of outstanding requests on a CIFS connection. |
| trusted_domain_enum_search_enabled | boolean | Specifies whether or not to enable trusted domain search.<br>• If this parameter is set to true, it displays CIFS options only for CIFS servers that support enumeration of bidirectional trusted domains and that support searching in all bidirectional trusted domains when performing Windows user lookups for UNIX user to Windows user name mapping.<br>• If set to false, it displays CIFS options for CIFS servers that do not support enumeration of bidirectional trusted domains. |
| widelink_reparse_versions | array[string] | Specifies the CIFS protocol versions for which the widelink is reported as reparse point. |

cifs_service_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |
| aes_netlogon_enabled | boolean | Specifies whether or not an AES session key is enabled for the Netlogon channel. |
| encrypt_dc_connection | boolean | Specifies whether encryption is required for domain controller connections. |

| Name | Type | Description |
| --- | --- | --- |
| kdc_encryption | boolean | This attribute has been deprecated. Use "security.advertised_kdc_encryptions" to specify the encryption type to use.<br>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.<br>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server. Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.<br>The CIFS server supports the following encryption types for Kerberos communication:<br>• RC4-HMAC<br>• DES<br>• AES<br>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the |

| Name | Type | Description |
| --- | --- | --- |
| ldap_referral_enabled | boolean | Specifies whether or not LDAP referral chasing is enabled for AD LDAP connections. |
| lm_compatibility_level | string | It is CIFS server minimum security level, also known as the LMCompatibilityLevel. The minimum security level is the minimum level of the security tokens that the CIFS server accepts from SMB clients.<br>The available values are:<br>• lm_ntlm_ntlmv2_krb Accepts LM, NTLM, NTLMv2 and Kerberos<br>• ntlm_ntlmv2_krb Accepts NTLM, NTLMv2 and Kerberos<br>• ntlmv2_krb Accepts NTLMv2 and Kerberos<br>• krb Accepts Kerberos only |
| restrict_anonymous | string | Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings.<br>The available values are:<br>• no_restriction - No access restriction for an anonymous user.<br>• no_enumeration - Enumeration is restricted for an anonymous user.<br>• no_access - All access is restricted for an anonymous user. |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies client session security for AD LDAP connections.<br>The available values are:<br>• none - No Signing or Sealing.<br>• sign - Sign LDAP traffic.<br>• seal - Seal and Sign LDAP traffic |
| smb_encryption | boolean | Specifies whether encryption is required for incoming CIFS traffic. |
| smb_signing | boolean | Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised. |
| try_ldap_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_ldaps | boolean | Specifies whether or not to use use LDAPS for secure Active Directory LDAP connections by using the TLS/SSL protocols. |
| use_start_tls | boolean | Specifies whether or not to use SSL/TLS for allowing secure LDAP communication with Active Directory LDAP servers. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_service

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ad_domain | ad_domain | |
| auth-style | string | Authentication type. |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |

| Name | Type | Description |
| --- | --- | --- |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |
| group_policy_object_enabled | boolean | If set to true, group policies will be applied to the SVM. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |

| Name | Type | Description |
| --- | --- | --- |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a CIFS server

POST `/protocols/cifs/services`

**Introduced In:** 9.6

Creates a CIFS server. Each SVM can have one CIFS server.

#### Important notes

- The CIFS server name might or might not be the same as the SVM name.
- The CIFS server name can contain up to 15 characters.
- The CIFS server name does not support the following characters: @ # * ( ) = + [ ] \| ; : " , < > / ?

#### Required properties when creating CIFS server with Windows Active Directory domain

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS server.
- `name` - Name of the CIFS server.
- `ad_domain.fqdn` - Fully qualified domain name of the Windows Active Directory to which this CIFS server belongs.
- `ad_domain.user` - User account with the access to add the CIFS server to the Active Directory.
- `ad_domain.password` - Account password used to add this CIFS server to the Active Directory.

#### Required properties when creating CIFS server in Workgroup mode

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS server.
- `name` - Name of the CIFS server.
- `workgroup` - Name of the workgroup to which this CIFS server belongs.

#### Required properties when using AKV for authentication (ANF platform)

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS server.
- `name` - Name of the CIFS server.
- `ad_domain.user` - User account with the access to add the CIFS server to the Active Directory.
- `ad_domain.fqdn` - Fully qualified domain name of the Windows Active Directory to which this CIFS server belongs.

- `client_id` - Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId.
- `tenant_id` - Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId.
- `key_vault_uri` - URI of the deployed AKV that is used by ONTAP for storing keys.
- `authentication_method` - Authentication method used by the application to prove its identity to AKV or EntraId. It can be either "client_secret" or "certificate".
- `auth_user_type` - Type of user who can access the SMB Volume. It can be either "domain_user" or "hybrid_user". The default is domain_user. In the case of a hybrid-user, ONTAP cannot access on-premise ADDS.
- `client_secret` - Secret used by the application to prove its identity to AKV.
- `client_certificate` - Base64 encoded PKCS12 certificate used by the application to prove its identity to AKV.

#### Recommended optional properties

- `comment` - Add a text comment of up to 256 characters about the CIFS server.
- `netbios.aliases` - Add a comma-delimited list of one or more NetBIOS aliases for the CIFS server.
- `netbios.wins_servers` - Add a list of Windows Internet Name Server (WINS) addresses that manage and map the NetBIOS name of the CIFS server to their network IP addresses. The IP addresses must be IPv4 addresses.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `ad_domain.organizational_unit` - *CN=Computers*
- `enabled` - *true*
- `security.restrict_anonymous` - *no_enumeration*
- `security.smb_signing` - *false*
- `security.smb_encryption` - *false*
- `security.encrypt_dc_connection` - *false*
- `default_unix_user` - *pcuser*
- `netbios.enabled` - *false* However, if either "netbios.wins-server" or "netbios.aliases" is set during POST and if `netbios.enabled` is not specified then `netbios.enabled` is set to true.
- `security.aes_netlogon_enabled` - *false*
- `security.try_ldap_channel_binding` - *true*
- `security.ldap_referral_enabled` - *false*

#### Related ONTAP commands

- `vserver cifs server create`
- `vserver cifs server options modify`

- `vserver cifs security modify`
- `vserver cifs server add-netbios-aliases`

**Learn more**

- DOC /protocols/cifs/services

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | If this is set and a machine account with the same name as specified in 'cifs-server name' exists in the Active Directory, existing machine account will be overwritten and reused.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| ad_domain | ad_domain | |
| auth-style | string | Authentication type. |

| Name | Type | Description |
| --- | --- | --- |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |

| Name | Type | Description |
| --- | --- | --- |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |
| proxy_password | string | Proxy password. Password is not audited. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |

| Name | Type | Description |
| --- | --- | --- |
| workgroup | string | The workgroup name. |

**Example request**

```
{
  "ad_domain": {
    "default_site": "string",
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string",
    "user": "string"
  },
  "auth-style": "domain",
  "auth_user_type": "string",
  "authentication_method": "string",
  "client_certificate": "PEM Cert",
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "client_secret": "<APPLICATION-CLIENT-SECRET-FOR-AKV>",
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "group_policy_object_enabled": true,
  "key_vault_uri": "https://kmip-akv-keyvault.vault.azure.net/",
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "oauth_host": "login.microsoftonline.com",
  "options": {
    "max_connections_per_session": 32,
    "max_lifs_per_session": 256,
    "max_opens_same_file_per_tree": 1000,
    "max_same_tree_connect_per_session": 5000,
    "max_same_user_sessions_per_connection": 2500,
    "max_watches_set_per_tree": 500,
    "null_user_windows_name": "string",
    "smb_credits": 128,
    "widelink_reparse_versions": [
      "smb1"
    ]
  },
```

```
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "string",
  "proxy_username": "proxyuser",
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ],
    "restrict_anonymous": "string",
    "session_security": "string"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
  "timeout": 25,
  "workgroup": "workgrp1"
}
```

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**

```
{
  "job": {
    "uuid": "string"
  }
}
```

**Headers**

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115127 | The cluster lacks a valid CIFS license. |
| 3735751 | Failed to authenticate and retrieve the access token from the Azure OAuth host. |
| 3735752 | Failed to extract the private key from the Azure Key Vault certificate. |
| 3735753 | Unsupported content_type in the Azure secrets response. |
| 3735754 | Failed to parse the JSON response from Azure Key Vault. |
| 3735755 | REST API call to Azure failed. |
| 3735756 | Invalid client certificate. |
| 3735757 | Failed to generate client assertion. |
| 3735762 | The provided Azure Key Vault configuration is incorrect. |
| 3735763 | The provided Azure Key Vault configuration is incomplete. |
| 3735764 | Request to Azure failed. Reason - Azure error code and Azure error message. |
| 655388 | Failed to create the CIFS server. Failure can occur due to multiple reasons. For example: STARTTLS and LDAPS cannot be used together or a password update failed or error from LDAP server or Kerberos Error: KDC has no support for encryption type. |
| 655524 | CIFS server creation failed. |
| 655538 | CIFS server creation failed because a server with the same name already exists. |
| 655562 | NetBIOS name is longer than 15 characters. |
| 655563 | NetBIOS name contains characters that are not allowed. |

| Error Code | Description |
| --- | --- |
| 655771 | The number of NetBIOS aliases cannot exceed the maximum supported number of '200'. |
| 655914 | Failed to create the Active Directory machine account. |
| 655923 | Retrieving credentials from AKV is not supported because the effective cluster version is not ONTAP 9.15.0 or later. |
| 656464 | Failed to create the Active Directory machine account. Reason: Invalid Credentials. |
| 656465 | Failed to create the Active Directory machine account. Reason: Account with same name already exists. |
| 656466 | Failed to create the Active Directory machine account. Reason: Domain Controller is not reachable or does not exist. |
| 656467 | Failed to create the Active Directory machine account. Reason: Organizational-Unit not found. |
| 656473 | Fields security.kdc_encryption and security.advertised_kdc_encryptions are mutually exclusive. Specify only one of the two. |
| 655477 | Only one CIFS server is supported per SVM. |
| 656487 | Missing fields when doing CIFS operation with hybrid auth user type. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only. |
| password | string | The account password used to add this CIFS server to the Active Directory. This is not audited. |
| user | string | The user account used to add this CIFS server to the Active Directory. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

cifs_netbios

| Name | Type | Description |
| --- | --- | --- |
| aliases | array[string] | |
| enabled | boolean | Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration. |
| wins_servers | array[string] | |

cifs_service_options

| Name | Type | Description |
| --- | --- | --- |
| admin_to_root_mapping | boolean | Specifies whether or not Administrator can be mapped to the UNIX user "root". |
| advanced_sparse_file | boolean | Specifies whether or not the CIFS server supports the advanced sparse file capabilities. This allows<br>CIFS clients to query the allocated ranges of a file and to write zeroes or free data blocks for ranges<br>of a file. |
| backup_symlink_enabled | boolean | Specifies whether or not to preserve UNIX symlinks during backup through SMB. |
| client_dup_detection_enabled | boolean | Specifies whether or not client duplicate session detection is enabled for CIFS. |
| client_version_reporting_enabled | boolean | Specifies whether or not client version reporting is enabled for CIFS. |

| Name | Type | Description |
| --- | --- | --- |
| copy_offload | boolean | Specifies whether or not to enable the Copy Offload feature. This feature enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.<br><br>Note that this will also enable/disable the direct copy feature accordingly. |
| dac_enabled | boolean | Specifies whether or not the Dynamic Access Control (DAC) feature is enabled for the CIFS server. |
| export_policy_enabled | boolean | Specifies whether or not export policies are enabled for CIFS. |
| fake_open | boolean | Specifies whether or not fake open support is enabled. This parameter allows you to optimize the<br>open and close requests coming from SMB 2 clients. |
| fsctl_trim | boolean | Specifies whether or not the trim requests (FSCTL_FILE_LEVEL_TRIM) are supported on the CIFS server. |
| junction_reparse | boolean | Specifies whether or not the reparse point support is enabled. When enabled the CIFS server exposes junction points to Windows clients as reparse points. This parameter is only active<br>if the client has negotiated use of the SMB 2 or SMB 3 protocol. This parameter is not supported for SVMs with Infinite Volume. |
| large_mtu | boolean | Specifies whether or not SMB clients can send reads up to 1 MB in size. |

| Name | Type | Description |
| --- | --- | --- |
| max_connections_per_session | integer | Specifies the maximum number of connections allowed per multichannel session. |
| max_lifs_per_session | integer | Specifies the maximum number of LIFs advertised per multichannel session. |
| max_opens_same_file_per_tree | integer | Specifies the maximum number of opens on the same file per tree. |
| max_same_tree_connect_per_session | integer | Specifies the maximum number of same tree connections per session. |
| max_same_user_sessions_per_connection | integer | Specifies the maximum number of same user sessions per connection. |
| max_watches_set_per_tree | integer | Specifies the maximum number of watches set per tree. |
| multichannel | boolean | Specifies whether or not the CIFS server supports Multichannel. |
| null_user_windows_name | string | Specifies a Windows User or Group name that should be mapped in case of a NULL user value. |
| path_component_cache | boolean | Specifies whether or not the path component cache is enabled on the CIFS server. |
| referral | boolean | Specifies whether or not to refer clients to more optimal LIFs. When enabled, it automatically refers clients to a data LIF local to the node which hosts the root of the requested share. |

| Name | Type | Description |
| --- | --- | --- |
| shadowcopy | boolean | Specifies whether or not to enable the Shadowcopy Feature. This feature enables to take share-based backup copies of data that is in a data-consistent state at a specific point in time where the data is accessed over SMB 3.0 shares. |
| shadowcopy_dir_depth | integer | Specifies the maximum level of subdirectories on which ONTAP should create shadow copies.<br>• Default value: 5<br>• Introduced in: 9.11<br>• x-nullable: true |
| smb_credits | integer | Specifies the maximum number of outstanding requests on a CIFS connection. |
| trusted_domain_enum_search_enabled | boolean | Specifies whether or not to enable trusted domain search.<br>• If this parameter is set to true, it displays CIFS options only for CIFS servers that support enumeration of bidirectional trusted domains and that support searching in all bidirectional trusted domains when performing Windows user lookups for UNIX user to Windows user name mapping.<br>• If set to false, it displays CIFS options for CIFS servers that do not support enumeration of bidirectional trusted domains. |
| widelink_reparse_versions | array[string] | Specifies the CIFS protocol versions for which the widelink is reported as reparse point. |

cifs_service_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |
| aes_netlogon_enabled | boolean | Specifies whether or not an AES session key is enabled for the Netlogon channel. |
| encrypt_dc_connection | boolean | Specifies whether encryption is required for domain controller connections. |

| Name | Type | Description |
| --- | --- | --- |
| kdc_encryption | boolean | This attribute has been deprecated. Use "security.advertised_kdc_encryptions" to specify the encryption type to use.<br>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.<br>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server. Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.<br>The CIFS server supports the following encryption types for Kerberos communication:<br>• RC4-HMAC<br>• DES<br>• AES<br>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the |

| Name | Type | Description |
| --- | --- | --- |
| ldap_referral_enabled | boolean | Specifies whether or not LDAP referral chasing is enabled for AD LDAP connections. |
| restrict_anonymous | string | Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names.<br>Access for the anonymous user can be controlled by specifying one of three access restriction settings.<br>The available values are:<br>• no_restriction - No access restriction for an anonymous user.<br>• no_enumeration - Enumeration is restricted for an anonymous user.<br>• no_access - All access is restricted for an anonymous user. |
| session_security | string | Specifies client session security for AD LDAP connections.<br>The available values are:<br>• none - No Signing or Sealing.<br>• sign - Sign LDAP traffic.<br>• seal - Seal and Sign LDAP traffic |
| smb_encryption | boolean | Specifies whether encryption is required for incoming CIFS traffic. |
| smb_signing | boolean | Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised. |

| Name | Type | Description |
| --- | --- | --- |
| try_ldap_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_ldaps | boolean | Specifies whether or not to use use LDAPS for secure Active Directory LDAP connections by using the TLS/SSL protocols. |
| use_start_tls | boolean | Specifies whether or not to use SSL/TLS for allowing secure LDAP communication with Active Directory LDAP servers. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_service

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | |
| auth-style | string | Authentication type. |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |
| proxy_password | string | Proxy password. Password is not audited. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |

| Name | Type | Description |
|---|---|---|
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete a CIFS server and related configurations**

DELETE `/protocols/cifs/services/{svm.uuid}`

**Introduced In:** 9.6

Deletes a CIFS server and related CIFS configurations.

Important notes:

- The default value for the "force" field is false.

- If the "force" field is set along with user login credentials, the local CIFS configuration will be deleted irrespective of any communication errors.
- If the "force" field alone is set without passing the user login credentials, the local CIFS configuration will be deleted by not making any request to Active Directory.

#### Related ONTAP commands

- `vserver cifs server delete`
- `vserver cifs remove-netbios-aliases`

#### Learn more

- DOC /protocols/cifs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | When set, the local CIFS configuration is deleted irrespective of any communication errors. Default value for this field is false.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| ad_domain | ad_domain_delete | |

| Name | Type | Description |
| --- | --- | --- |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| oauth_host | string | Open authorization server host name. |
| proxy_host | string | Proxy host. |
| proxy_password | string | Proxy password. Password is not audited. |

| Name | Type | Description |
| --- | --- | --- |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

#### Example request

```
{
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string",
    "user": "string"
  },
  "auth_user_type": "string",
  "authentication_method": "string",
  "client_certificate": "PEM Cert",
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "client_secret": "<APPLICATION-CLIENT-SECRET-FOR-AKV>",
  "key_vault_uri": "https://kmip-akv-keyvault.vault.azure.net/",
  "oauth_host": "login.microsoftonline.com",
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "string",
  "proxy_username": "proxyuser",
  "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
  "timeout": 25,
  "workgroup": "workgrp1"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655525 | In order to delete an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges. |
| 3735751 | Failed to authenticate and retrieve the access token from the Azure OAuth host. |
| 3735752 | Failed to extract the private key from the Azure Key Vault certificate. |
| 3735753 | Unsupported content_type in the Azure secrets response. |
| 3735754 | Failed to parse the JSON response from Azure Key Vault. |
| 3735755 | REST API call to Azure failed. |
| 3735756 | Invalid client certificate. |
| 3735757 | Failed to generate client assertion. |

| Error Code | Description |
| --- | --- |
| 3735762 | The provided Azure Key Vault configuration is incorrect. |
| 3735763 | The provided Azure Key Vault configuration is incomplete. |
| 3735764 | Request to Azure failed. Reason - Azure error code and Azure error message. |
| 655563 | NetBIOS name contains characters that are not allowed. |
| 655562 | NetBIOS name is longer than 15 characters. |
| 655923 | Retrieving credentials from AKV is not supported because the effective cluster version is not ONTAP 9.15.0 or later. |
| 656487 | Missing fields when doing CIFS operation with hybrid auth user type. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

ad_domain_delete

| Name | Type | Description |
|---|---|---|
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only. |
| password | string | The account password used to add this CIFS server to the Active Directory. This is not audited. |
| user | string | The user account used to add this CIFS server to the Active Directory. |

cifs_service_delete

| Name | Type | Description |
|---|---|---|
| ad_domain | ad_domain_delete | |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| oauth_host | string | Open authorization server host name. |
| proxy_host | string | Proxy host. |
| proxy_password | string | Proxy password. Password is not audited. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |

| Name | Type | Description |
| --- | --- | --- |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a CIFS server

GET `/protocols/cifs/services/{svm.uuid}`

**Introduced In:** 9.6

Retrieves a CIFS server.

#### Related ONTAP commands

- `vserver cifs server show`
- `vserver cifs server options show`
- `vserver cifs server security show`

#### Learn more

- DOC /protocols/cifs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ad_domain | ad_domain | |

| Name | Type | Description |
| --- | --- | --- |
| auth-style | string | Authentication type. |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br><br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br><br>• client_secret<br>• certificate |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |
| group_policy_object_enabled | boolean | If set to true, group policies will be applied to the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |

| Name | Type | Description |
| --- | --- | --- |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "default_site": "string",
    "fqdn": "example.com",
    "organizational_unit": "string"
  },
  "auth-style": "domain",
  "auth_user_type": "string",
  "authentication_method": "string",
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "key_vault_uri": "https://kmip-akv-keyvault.vault.azure.net/",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "CIFS1",
```

```
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "oauth_host": "login.microsoftonline.com",
  "options": {
    "max_connections_per_session": 32,
    "max_lifs_per_session": 256,
    "max_opens_same_file_per_tree": 1000,
    "max_same_tree_connect_per_session": 5000,
    "max_same_user_sessions_per_connection": 2500,
    "max_watches_set_per_tree": 500,
    "null_user_windows_name": "string",
    "smb_credits": 128,
    "widelink_reparse_versions": [
      "smb1"
    ]
  },
  "proxy_host": "proxy.eng.com",
  "proxy_port": 1234,
  "proxy_type": "string",
  "proxy_username": "proxyuser",
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ],
    "lm_compatibility_level": "string",
    "restrict_anonymous": "string",
    "session_security": "string"
  },
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
```

```
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
  "timeout": 25,
  "workgroup": "workgrp1"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

cifs_netbios

| Name | Type | Description |
| --- | --- | --- |
| aliases | array[string] | |
| enabled | boolean | Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration. |
| wins_servers | array[string] | |

cifs_service_options

| Name | Type | Description |
| --- | --- | --- |
| admin_to_root_mapping | boolean | Specifies whether or not Administrator can be mapped to the UNIX user "root". |
| advanced_sparse_file | boolean | Specifies whether or not the CIFS server supports the advanced sparse file capabilities. This allows<br>CIFS clients to query the allocated ranges of a file and to write zeroes or free data blocks for ranges<br>of a file. |
| backup_symlink_enabled | boolean | Specifies whether or not to preserve UNIX symlinks during backup through SMB. |
| client_dup_detection_enabled | boolean | Specifies whether or not client duplicate session detection is enabled for CIFS. |
| client_version_reporting_enabled | boolean | Specifies whether or not client version reporting is enabled for CIFS. |

| Name | Type | Description |
| --- | --- | --- |
| copy_offload | boolean | Specifies whether or not to enable the Copy Offload feature. This feature enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.<br><br>Note that this will also enable/disable the direct copy feature accordingly. |
| dac_enabled | boolean | Specifies whether or not the Dynamic Access Control (DAC) feature is enabled for the CIFS server. |
| export_policy_enabled | boolean | Specifies whether or not export policies are enabled for CIFS. |
| fake_open | boolean | Specifies whether or not fake open support is enabled. This parameter allows you to optimize the<br>open and close requests coming from SMB 2 clients. |
| fsctl_trim | boolean | Specifies whether or not the trim requests (FSCTL_FILE_LEVEL_TRIM) are supported on the CIFS server. |
| junction_reparse | boolean | Specifies whether or not the reparse point support is enabled. When enabled the CIFS server exposes junction points to Windows clients as reparse points. This parameter is only active<br>if the client has negotiated use of the SMB 2 or SMB 3 protocol. This parameter is not supported for SVMs with Infinite Volume. |
| large_mtu | boolean | Specifies whether or not SMB clients can send reads up to 1 MB in size. |

| Name | Type | Description |
| --- | --- | --- |
| max_connections_per_session | integer | Specifies the maximum number of connections allowed per multichannel session. |
| max_lifs_per_session | integer | Specifies the maximum number of LIFs advertised per multichannel session. |
| max_opens_same_file_per_tree | integer | Specifies the maximum number of opens on the same file per tree. |
| max_same_tree_connect_per_session | integer | Specifies the maximum number of same tree connections per session. |
| max_same_user_sessions_per_connection | integer | Specifies the maximum number of same user sessions per connection. |
| max_watches_set_per_tree | integer | Specifies the maximum number of watches set per tree. |
| multichannel | boolean | Specifies whether or not the CIFS server supports Multichannel. |
| null_user_windows_name | string | Specifies a Windows User or Group name that should be mapped in case of a NULL user value. |
| path_component_cache | boolean | Specifies whether or not the path component cache is enabled on the CIFS server. |
| referral | boolean | Specifies whether or not to refer clients to more optimal LIFs. When enabled, it automatically refers clients to a data LIF local to the node which hosts the root of the requested share. |

| Name | Type | Description |
| --- | --- | --- |
| shadowcopy | boolean | Specifies whether or not to enable the Shadowcopy Feature. This feature enables to take share-based backup copies of data that is in a data-consistent state at a specific point in time where the data is accessed over SMB 3.0 shares. |
| shadowcopy_dir_depth | integer | Specifies the maximum level of subdirectories on which ONTAP should create shadow copies.<br>• Default value: 5<br>• Introduced in: 9.11<br>• x-nullable: true |
| smb_credits | integer | Specifies the maximum number of outstanding requests on a CIFS connection. |
| trusted_domain_enum_search_enabled | boolean | Specifies whether or not to enable trusted domain search.<br>• If this parameter is set to true, it displays CIFS options only for CIFS servers that support enumeration of bidirectional trusted domains and that support searching in all bidirectional trusted domains when performing Windows user lookups for UNIX user to Windows user name mapping.<br>• If set to false, it displays CIFS options for CIFS servers that do not support enumeration of bidirectional trusted domains. |
| widelink_reparse_versions | array[string] | Specifies the CIFS protocol versions for which the widelink is reported as reparse point. |

cifs_service_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |
| aes_netlogon_enabled | boolean | Specifies whether or not an AES session key is enabled for the Netlogon channel. |
| encrypt_dc_connection | boolean | Specifies whether encryption is required for domain controller connections. |

| Name | Type | Description |
| --- | --- | --- |
| kdc_encryption | boolean | This attribute has been deprecated. Use "security.advertised_kdc_encryptions" to specify the encryption type to use.<br>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.<br>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server. Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.<br>The CIFS server supports the following encryption types for Kerberos communication:<br>• RC4-HMAC<br>• DES<br>• AES<br>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the |

| Name | Type | Description |
| --- | --- | --- |
| ldap_referral_enabled | boolean | Specifies whether or not LDAP referral chasing is enabled for AD LDAP connections. |
| lm_compatibility_level | string | It is CIFS server minimum security level, also known as the LMCompatibilityLevel. The minimum security level is the minimum level of the security tokens that the CIFS server accepts from SMB clients.<br>The available values are:<br>• lm_ntlm_ntlmv2_krb Accepts LM, NTLM, NTLMv2 and Kerberos<br>• ntlm_ntlmv2_krb Accepts NTLM, NTLMv2 and Kerberos<br>• ntlmv2_krb Accepts NTLMv2 and Kerberos<br>• krb Accepts Kerberos only |
| restrict_anonymous | string | Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings.<br>The available values are:<br>• no_restriction - No access restriction for an anonymous user.<br>• no_enumeration - Enumeration is restricted for an anonymous user.<br>• no_access - All access is restricted for an anonymous user. |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies client session security for AD LDAP connections.<br>The available values are:<br>• none - No Signing or Sealing.<br>• sign - Sign LDAP traffic.<br>• seal - Seal and Sign LDAP traffic |
| smb_encryption | boolean | Specifies whether encryption is required for incoming CIFS traffic. |
| smb_signing | boolean | Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised. |
| try_ldap_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_ldaps | boolean | Specifies whether or not to use use LDAPS for secure Active Directory LDAP connections by using the TLS/SSL protocols. |
| use_start_tls | boolean | Specifies whether or not to use SSL/TLS for allowing secure LDAP communication with Active Directory LDAP servers. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update CIFS mandatory and optional parameters**

PATCH `/protocols/cifs/services/{svm.uuid}`

**Introduced In:** 9.6

Updates both the mandatory and optional parameters of the CIFS configuration. Ensure the CIFS server is administratively disabled when renaming the CIFS server or modifying the *ad_domain* properties.

#### Related ONTAP commands

- `vserver cifs server modify`
- `vserver cifs server options modify`
- `vserver cifs security modify`
- `vserver cifs server add-netbios-aliases`
- `vserver cifs server remove-netbios-aliases`
- `vserver cifs group-policy modify`

#### Learn more

- DOC /protocols/cifs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | If this is set and a machine account with the same name as specified in 'cifs-server name' exists in the Active Directory, existing machine account will be overwritten and reused. The default value for this field is false.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| ad_domain | ad_domain | |
| auth-style | string | Authentication type. |

| Name | Type | Description |
| --- | --- | --- |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |

| Name | Type | Description |
| --- | --- | --- |
| group_policy_object_enabled | boolean | If set to true, group policies will be applied to the SVM. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |
| proxy_password | string | Proxy password. Password is not audited. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |

| Name | Type | Description |
| --- | --- | --- |
| workgroup | string | The workgroup name. |

**Example request**

```
{
  "ad_domain": {
    "default_site": "string",
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string",
    "user": "string"
  },
  "auth-style": "domain",
  "auth_user_type": "string",
  "authentication_method": "string",
  "client_certificate": "PEM Cert",
  "client_id": "e959d1b5-5a63-4284-9268-851e30e3eceb",
  "client_secret": "<APPLICATION-CLIENT-SECRET-FOR-AKV>",
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "key_vault_uri": "https://kmip-akv-keyvault.vault.azure.net/",
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "oauth_host": "login.microsoftonline.com",
  "options": {
    "max_connections_per_session": 32,
    "max_lifs_per_session": 256,
    "max_opens_same_file_per_tree": 1000,
    "max_same_tree_connect_per_session": 5000,
    "max_same_user_sessions_per_connection": 2500,
    "max_watches_set_per_tree": 500,
    "null_user_windows_name": "string",
    "smb_credits": 128,
    "widelink_reparse_versions": [
      "smb1"
    ]
  },
  "proxy_host": "proxy.eng.com",
```

```
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "string",
  "proxy_username": "proxyuser",
  "security": {
    "advertised_kdc_encryptions": [
      "string"
    ],
    "lm_compatibility_level": "string",
    "restrict_anonymous": "string",
    "session_security": "string"
  },
  "tenant_id": "c9f32fcb-4ab7-40fe-af1b-1850d46cfbbe",
  "timeout": 25,
  "workgroup": "workgrp1"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**

```
{
  "job": {
    "uuid": "string"
  }
}
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 3735751 | Failed to authenticate and retrieve the access token from the Azure OAuth host. |
| 3735752 | Failed to extract the private key from the Azure Key Vault certificate. |
| 3735753 | Unsupported content_type in the Azure secrets response. |
| 3735754 | Failed to parse the JSON response from Azure Key Vault. |
| 3735755 | REST API call to Azure failed. |
| 3735756 | Invalid client certificate. |
| 3735757 | Failed to generate client assertion. |
| 3735762 | The provided Azure Key Vault configuration is incorrect. |
| 3735763 | The provided Azure Key Vault configuration is incomplete. |
| 3735764 | Request to Azure failed. Reason - Azure error code and Azure error message. |
| 655390 | Failed to modify the CIFS server. Failure can occur due to multiple reasons. For example: STARTTLS and LDAPS cannot be used together or the group policy is already enabled for the SVM. |
| 655562 | NetBIOS name is longer than 15 characters. |
| 655538 | Unable to modify the CIFS server. The server name is already used by another SVM. |
| 655563 | NetBIOS name contains characters that are not allowed. |
| 655771 | The number of NetBIOS aliases cannot exceed the maximum supported number of '200'. |
| 655923 | Retrieving credentials from AKV is not supported because the effective cluster version is not ONTAP 9.15.0 or later. |
| 656473 | Fields security.kdc_encryption and security.advertised_kdc_encryptions are mutually exclusive. Specify only one of the two. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only. |
| password | string | The account password used to add this CIFS server to the Active Directory. This is not audited. |
| user | string | The user account used to add this CIFS server to the Active Directory. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

cifs_netbios

| Name | Type | Description |
| --- | --- | --- |
| aliases | array[string] | |
| enabled | boolean | Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration. |
| wins_servers | array[string] | |

cifs_service_options

| Name | Type | Description |
| --- | --- | --- |
| admin_to_root_mapping | boolean | Specifies whether or not Administrator can be mapped to the UNIX user "root". |
| advanced_sparse_file | boolean | Specifies whether or not the CIFS server supports the advanced sparse file capabilities. This allows<br>CIFS clients to query the allocated ranges of a file and to write zeroes or free data blocks for ranges<br>of a file. |
| backup_symlink_enabled | boolean | Specifies whether or not to preserve UNIX symlinks during backup through SMB. |
| client_dup_detection_enabled | boolean | Specifies whether or not client duplicate session detection is enabled for CIFS. |
| client_version_reporting_enabled | boolean | Specifies whether or not client version reporting is enabled for CIFS. |

| Name | Type | Description |
| --- | --- | --- |
| copy_offload | boolean | Specifies whether or not to enable the Copy Offload feature. This feature enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.<br><br>Note that this will also enable/disable the direct copy feature accordingly. |
| dac_enabled | boolean | Specifies whether or not the Dynamic Access Control (DAC) feature is enabled for the CIFS server. |
| export_policy_enabled | boolean | Specifies whether or not export policies are enabled for CIFS. |
| fake_open | boolean | Specifies whether or not fake open support is enabled. This parameter allows you to optimize the<br>open and close requests coming from SMB 2 clients. |
| fsctl_trim | boolean | Specifies whether or not the trim requests (FSCTL_FILE_LEVEL_TRIM) are supported on the CIFS server. |
| junction_reparse | boolean | Specifies whether or not the reparse point support is enabled. When enabled the CIFS server exposes junction points to Windows clients as reparse points. This parameter is only active<br>if the client has negotiated use of the SMB 2 or SMB 3 protocol. This parameter is not supported for SVMs with Infinite Volume. |
| large_mtu | boolean | Specifies whether or not SMB clients can send reads up to 1 MB in size. |

| Name | Type | Description |
| --- | --- | --- |
| max_connections_per_session | integer | Specifies the maximum number of connections allowed per multichannel session. |
| max_lifs_per_session | integer | Specifies the maximum number of LIFs advertised per multichannel session. |
| max_opens_same_file_per_tree | integer | Specifies the maximum number of opens on the same file per tree. |
| max_same_tree_connect_per_session | integer | Specifies the maximum number of same tree connections per session. |
| max_same_user_sessions_per_connection | integer | Specifies the maximum number of same user sessions per connection. |
| max_watches_set_per_tree | integer | Specifies the maximum number of watches set per tree. |
| multichannel | boolean | Specifies whether or not the CIFS server supports Multichannel. |
| null_user_windows_name | string | Specifies a Windows User or Group name that should be mapped in case of a NULL user value. |
| path_component_cache | boolean | Specifies whether or not the path component cache is enabled on the CIFS server. |
| referral | boolean | Specifies whether or not to refer clients to more optimal LIFs. When enabled, it automatically refers clients to a data LIF local to the node which hosts the root of the requested share. |

| Name | Type | Description |
| --- | --- | --- |
| shadowcopy | boolean | Specifies whether or not to enable the Shadowcopy Feature. This feature enables to take share-based backup copies of data that is in a data-consistent state at a specific point in time where the data is accessed over SMB 3.0 shares. |
| shadowcopy_dir_depth | integer | Specifies the maximum level of subdirectories on which ONTAP should create shadow copies.<br>• Default value: 5<br>• Introduced in: 9.11<br>• x-nullable: true |
| smb_credits | integer | Specifies the maximum number of outstanding requests on a CIFS connection. |
| trusted_domain_enum_search_enabled | boolean | Specifies whether or not to enable trusted domain search.<br>• If this parameter is set to true, it displays CIFS options only for CIFS servers that support enumeration of bidirectional trusted domains and that support searching in all bidirectional trusted domains when performing Windows user lookups for UNIX user to Windows user name mapping.<br>• If set to false, it displays CIFS options for CIFS servers that do not support enumeration of bidirectional trusted domains. |
| widelink_reparse_versions | array[string] | Specifies the CIFS protocol versions for which the widelink is reported as reparse point. |

cifs_service_security

| Name | Type | Description |
| --- | --- | --- |
| advertised_kdc_encryptions | array[string] | |
| aes_netlogon_enabled | boolean | Specifies whether or not an AES session key is enabled for the Netlogon channel. |
| encrypt_dc_connection | boolean | Specifies whether encryption is required for domain controller connections. |

| Name | Type | Description |
| --- | --- | --- |
| kdc_encryption | boolean | This attribute has been deprecated. Use "security.advertised_kdc_encryptions" to specify the encryption type to use.<br>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.<br>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server. Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.<br>The CIFS server supports the following encryption types for Kerberos communication:<br>• RC4-HMAC<br>• DES<br>• AES<br>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the |

| Name | Type | Description |
| --- | --- | --- |
| ldap_referral_enabled | boolean | Specifies whether or not LDAP referral chasing is enabled for AD LDAP connections. |
| lm_compatibility_level | string | It is CIFS server minimum security level, also known as the LMCompatibilityLevel. The minimum security level is the minimum level of the security tokens that the CIFS server accepts from SMB clients.<br>The available values are:<br>• lm_ntlm_ntlmv2_krb Accepts LM, NTLM, NTLMv2 and Kerberos<br>• ntlm_ntlmv2_krb Accepts NTLM, NTLMv2 and Kerberos<br>• ntlmv2_krb Accepts NTLMv2 and Kerberos<br>• krb Accepts Kerberos only |
| restrict_anonymous | string | Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings.<br>The available values are:<br>• no_restriction - No access restriction for an anonymous user.<br>• no_enumeration - Enumeration is restricted for an anonymous user.<br>• no_access - All access is restricted for an anonymous user. |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies client session security for AD LDAP connections.<br>The available values are:<br>• none - No Signing or Sealing.<br>• sign - Sign LDAP traffic.<br>• seal - Seal and Sign LDAP traffic |
| smb_encryption | boolean | Specifies whether encryption is required for incoming CIFS traffic. |
| smb_signing | boolean | Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised. |
| try_ldap_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_ldaps | boolean | Specifies whether or not to use use LDAPS for secure Active Directory LDAP connections by using the TLS/SSL protocols. |
| use_start_tls | boolean | Specifies whether or not to use SSL/TLS for allowing secure LDAP communication with Active Directory LDAP servers. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_service

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | |
| auth-style | string | Authentication type. |
| auth_user_type | string | Specifies the type of user who can access the SMB Volume. The default is domain_user. In the case of a hybrid-user, ONTAP won't contact on-premise ADDS.<br>• enum: ["domain_user", "hybrid_user"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| authentication_method | string | Specifies the authentication method.<br>The available values are:<br>• client_secret<br>• certificate |
| client_certificate | string | PKCS12 certificate used by the application to prove its identity to AKV. |

| Name | Type | Description |
| --- | --- | --- |
| client_id | string | Application client ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| client_secret | string | Secret used by the application to prove its identity to AKV. |
| comment | string | A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks. |
| default_unix_user | string | Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails. |
| enabled | boolean | Specifies if the CIFS service is administratively enabled. |
| group_policy_object_enabled | boolean | If set to true, group policies will be applied to the SVM. |
| key_vault_uri | string | URI of the deployed AKV that is used by ONTAP for storing keys.<br>• example: https://kmip-akv-keyvault.vault.azure.net/<br>• format: uri<br>• Introduced in: 9.15<br>• x-nullable: true |
| name | string | The name of the CIFS server. |
| netbios | cifs_netbios | |
| oauth_host | string | Open authorization server host name. |
| options | cifs_service_options | |
| proxy_host | string | Proxy host. |

| Name | Type | Description |
| --- | --- | --- |
| proxy_password | string | Proxy password. Password is not audited. |
| proxy_port | integer | Proxy port. |
| proxy_type | string | Proxy type. |
| proxy_username | string | Proxy username. |
| security | cifs_service_security | |
| tenant_id | string | Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV or EntraId. |
| timeout | integer | AKV connection timeout, in seconds. The allowed range is between 0 to 30 seconds. |
| verify_host | boolean | Verify the identity of the AKV host name. By default, verify_host is set to true. |
| workgroup | string | The workgroup name. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve CIFS protocol historical performance metrics for an SVM

GET `/protocols/cifs/services/{svm.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for the CIFS protocol of an SVM.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| duration | string | query | False | Filter by duration |
| status | string | query | False | Filter by status |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| throughput.other | integer | query | False | Filter by throughput.other |
| throughput.total | integer | query | False | Filter by throughput.total |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| timestamp | string | query | False | Filter by timestamp |
| svm.uuid | string | path | True | Unique identifier of the SVM. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View CIFS active file sessions

### View CIFS active file sessions

#### Overview

ONTAP `cifs sessions file show` functionality is used to provide a list of currently opened files.<h2/>

### Information on the open files

- Lists all files opened in current session.

#### Example

### Retrieving established open file information

To retrieve the list of open files, use the following API. Note that *return_records=true*.

```
# The API:
GET /protocols/cifs/session/files

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/session/files?fields=*&return_timeout=15&return_rec
ords=true" -H "accept: application/json"
```

```
# The response:
{
"records": [
                {
                            "node": {
                              "uuid": "a5f65ec0-3550-11ee-93c5-
005056ae78de",
                              "name": "sti220-vsim-sr050u"
                            },
                            "svm": {
                              "name": "vs0",
                              "uuid": "80e795f4-3553-11ee-9f97-
005056ae78de"
                            },
                            "identifier": 109,
                            "connection": {
                              "identifier": 103985,
                              "count": 1
                            },
                            "session": {
                              "identifier": 10878444899913433090
                            },
                            "type": "regular",
                            "open_mode": "r",
                            "volume": {
                              "uuid": "8384f6ae-3553-11ee-a3c3-
005056ae0dd5",
                              "name": "root_vs0"
                            },
                            "share": {
                              "name": "sh1",
                              "mode": "r"
                            },
                            "path": "first_file.txt",
                            "continuously_available": "no",
                            "range_locks_count": 0
                },
                {
                            "node": {
                              "uuid": "a5f65ec0-3550-11ee-93c5-
005056ae78de",
                              "name": "sti220-vsim-sr050u"
                            },
                            "svm": {
                              "name": "vs0",
                              "uuid": "80e795f4-3553-11ee-9f97-
```

```
005056ae78de"
                            },
                            "identifier": 110,
                            "connection": {
                              "identifier": 103985,
                              "count": 1
                            },
                            "session": {
                              "identifier": 10878444899913433090
                            },
                            "type": "regular",
                            "open_mode": "r",
                            "volume": {
                              "uuid": "8384f6ae-3553-11ee-a3c3-
005056ae0dd5",
                              "name": "root_vs0"
                             },
                             "share": {
                               "name": "sh1",
                               "mode": "r"
                             },
                             "path": "second_file.txt",
                             "continuously_available": "no",
                             "range_locks_count": 0
                }
           ],
           "num_records": 2
}
```


### Retrieving specific open file Information

```
# The API:
GET
/protocols/cifs/session/files/{node.uuid}/{svm.uuid}/{file.identifier}/{co
nnection.identifier}/{session_id}

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/session/files/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-
005056ae78de/109/103985/10878444899913433090"

# The response:
        {
              "node": {
                "uuid": "a5f65ec0-3550-11ee-93c5-005056ae78de",
                "name": "sti220-vsim-sr050u"
              },
              "svm": {
                "uuid": "80e795f4-3553-11ee-9f97-005056ae78de",
                "name": "vs0"
              },
              "identifier": 109,
              "connection": {
                "identifier": 103985,
                "count": 1
              },
              "session": {
                "identifier": 10878444899913433000
              },
              "type": "regular",
              "open_mode": "r",
              "volume": {
                "uuid": "8384f6ae-3553-11ee-a3c3-005056ae0dd5",
                "name": "root_vs0"
              },
              "share": {
                "name": "sh1",
                "mode": "r"
              },
              "path": "first_file.txt",
              "continuously_available": "no",
              "range_locks_count": 0
        }
```

**Closing a specific file based on `file.identifier`, `connection.identifier` and `session_id`**

The file being closed is identified by the UUID of its SVM, the corresponding file.identifier, connection.identifier and session_id.

```
# The API:
DELETE
/protocols/cifs/session/files/{node.uuid}/{svm.uuid}/{file.identifier}/{co
nnection.identifier}/{session_id}

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/protocols/cifs/session/files/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-
005056ae78de/109/103985/10878444899913433090" -H "accept:
application/json"
```

---

**Retrieve CIFS open files**

GET `/protocols/cifs/session/files`

**Introduced In:** 9.11

Retrieves the CIFS sessions Open Files information for all SVMs.

**Related ONTAP commands**

- `vserver cifs session file show`

**Learn more**

- DOC /protocols/cifs/session/files

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| continuously_available | string | query | False | Filter by continuously_available |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share.name | string | query | False | Filter by share.name |
| share.mode | string | query | False | Filter by share.mode |
| path | string | query | False | Filter by path |
| identifier | integer | query | False | Filter by identifier |
| type | string | query | False | Filter by type |
| open_mode | string | query | False | Filter by open_mode |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| connection.identifier | integer | query | False | Filter by connection.identifier |
| connection.count | integer | query | False | Filter by connection.count |
| range_locks_count | integer | query | False | Filter by range_locks_count |
| session.identifier | integer | query | False | Filter by session.identifier |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[cifs_open_file] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "connection": {
        "count": 3,
        "identifier": 356756
      },
      "continuously_available": "no",
      "identifier": 17,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "open_mode": "r",
      "path": "sub1\\sub2\\f4",
      "range_locks_count": 4,
      "session": {
        "identifier": 8966666858094657537
      },
      "share": {
        "mode": "r",
        "name": "share1"
      },
      "svm": {
        "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "type": "regular",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

connection

| Name | Type | Description |
|---|---|---|
| count | integer | The number of CIFS connections associated with the CIFS session. |
| identifier | integer | The connection that is used to open the file. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

session

| Name | Type | Description |
| --- | --- | --- |
| identifier | integer | Session under which file is opened. |

share

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The share mode used to open the file.<br>The share mode can be a combination of:<br>• r: read mode<br>• w: write mode<br>• d: delete<br>• rw: read and write modes<br>• rwd: read, write, and delete modes |
| name | string | CIFS share name where the file resides. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

cifs_open_file

Manage opened files over CIFS.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| connection | connection | |
| continuously_available | string | The type of continuous availability protection provided to the file. Opened files are continuously available if there are opened through a SMB3 client through a share with "continuously_available" set to yes.<br>These open files are capable of non-disruptively recovering from take over and giveback as well as general aggregate relocation.<br>• no: the open file is not continuously available.<br>• yes: the open file is continuously available. |
| identifier | integer | The unique identifier for the opened file. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| open_mode | string | Open mode corresponding to the opened file<br>• r: Opened for read<br>• w: Opened for write<br>• d: Opened for delete<br>• rw: Opened for read and write |
| path | string | Path from CIFS share. |
| range_locks_count | integer | The number of range locks granted on the file. |
| session | session | |
| share | share | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | The type of opened file.<br>The file can be a regular file, directory, a symbolic link to an existing file/directory, or an alternate data stream. |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Close open files identified by SVM

DELETE `/protocols/cifs/session/files/{node.uuid}/{svm.uuid}/{identifier}/{connection.identifier}/{session.identifier}`

**Introduced In:** 9.11

Closes open files identified by svm.uuid, file.identifier, connection.identifier and session_id.

#### Related ONTAP commands

- `vserver cifs session file close`

#### Learn more

- DOC /protocols/cifs/session/files

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| node.uuid | string | path | True | Node UUID. |
| identifier | integer | path | True | File ID<br>• format: int64 |
| connection.identifier | integer | path | True | Connection ID |
| session.identifier | integer | path | True | Session ID<br>• format: int64 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a CIFS open file

GET `/protocols/cifs/session/files/{node.uuid}/{svm.uuid}/{identifier}/{connection.identifier}/{session.identifier}`

**Introduced In:** 9.11

Retrieves specific CIFS session Open Files information.

#### Learn more

- DOC /protocols/cifs/session/files

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID. |
| identifier | integer | path | True | File ID<br>• format: int64 |
| connection.identifier | integer | path | True | Connection ID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| session.identifier | integer | path | True | Session ID<br>• format: int64 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| connection | connection | |
| continuously_available | string | The type of continuous availability protection provided to the file. Opened files are continuously available if there are opened through a SMB3 client through a share with "continuously_available" set to yes.<br>These open files are capable of non-disruptively recovering from take over and giveback as well as general aggregate relocation.<br>• no: the open file is not continuously available.<br>• yes: the open file is continuously available. |
| identifier | integer | The unique identifier for the opened file. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| open_mode | string | Open mode corresponding to the opened file<br>• r: Opened for read<br>• w: Opened for write<br>• d: Opened for delete<br>• rw: Opened for read and write |
| path | string | Path from CIFS share. |
| range_locks_count | integer | The number of range locks granted on the file. |
| session | session | |
| share | share | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | The type of opened file.<br>The file can be a regular file, directory, a symbolic link to an existing file/directory, or an alternate data stream. |
| volume | volume | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "connection": {
    "count": 3,
    "identifier": 356756
  },
  "continuously_available": "no",
  "identifier": 17,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "open_mode": "r",
  "path": "sub1\\sub2\\f4",
  "range_locks_count": 4,
  "session": {
    "identifier": 8966666858094657537
  },
  "share": {
    "mode": "r",
    "name": "share1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "regular",
  "volume": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

connection

| Name | Type | Description |
|---|---|---|
| count | integer | The number of CIFS connections associated with the CIFS session. |
| identifier | integer | The connection that is used to open the file. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

session

| Name | Type | Description |
|---|---|---|
| identifier | integer | Session under which file is opened. |

share

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The share mode used to open the file.<br>The share mode can be a combination of:<br>• r: read mode<br>• w: write mode<br>• d: delete<br>• rw: read and write modes<br>• rwd: read, write, and delete modes |
| name | string | CIFS share name where the file resides. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View CIFS active sessions

### View CIFS active sessions

#### Overview

ONTAP CIFS sessions `show` functionality is used to provide a list of currently established CIFS sessions with SMB Clients.

### Information on the CIFS session

- List all the SMB sessions for SVM and the clients along with volume information on which the clients are using.

#### Example

### Retrieves established sessions information

To retrieve the list of CIFS sessions, use the following API. Note that *return_records=true*.

```
# The API:
GET /protocols/cifs/sessions

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions?fields=*&return_timeout=15&return_records=
true" -H "accept: application/json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "a5f65ec0-3550-11ee-93c5-005056ae78de",
      "name": "sti220-vsim-sr050u"
    },
    "svm": {
      "name": "vs0",
      "uuid": "80e795f4-3553-11ee-9f97-005056ae78de"
    },
    "identifier": 10878444899913433000,
    "connection_id": 103985,
    "server_ip": "172.29.212.166",
    "client_ip": "10.249.30.206",
    "authentication": "ntlmv2",
    "user": "example_user",
    "mapped_unix_user": "example_user",
    "open_shares": 1,
    "open_files": 5,
    "open_other": 0,
    "connected_duration": "PT37M53S",
    "idle_duration": "PT3M15S",
    "protocol": "smb3_1",
    "continuous_availability": "unavailable",
    "smb_signing": true,
    "smb_encryption": "unencrypted",
    "connection_count": 1,
    "large_mtu": true,
    "volumes": [
```

```
      {
        "name": "root_vs0",
        "uuid": "5015c99b-3554-11ee-9f97-005056ae78de"
      },
      {
        "name": "vol_test",
        "uuid": "84b94bb8-3553-11ee-a3c3-005056ae0dd5"
      }
    ]
  },
  {
    "node": {
      "uuid": "a5f65ec0-3550-11ee-93c5-005056ae78de",
      "name": "sti220-vsim-sr050u"
    },
    "svm": {
      "name": "vs0",
      "uuid": "80e795f4-3553-11ee-9f97-005056ae78de"
    },
    "identifier": 10878444899913433000,
    "connection_id": 103986,
    "server_ip": "172.29.212.166",
    "client_ip": "10.234.157.109",
    "authentication": "ntlmv2",
    "user": "example_user",
    "mapped_unix_user": "example_user",
    "open_shares": 1,
    "open_files": 1,
    "open_other": 0,
    "connected_duration": "PT15S",
    "idle_duration": "PT5S",
    "protocol": "smb3_1",
    "continuous_availability": "unavailable",
    "smb_signing": false,
    "smb_encryption": "unencrypted",
    "connection_count": 1,
    "large_mtu": true,
    "volumes": [
      {
        "name": "root_vs0",
        "uuid": "84b94bb8-3553-11ee-a3c3-005056ae0dd5"
      }
    ]
  }
],
"num_records": 2
```

```
}
```

### Retrieving CIFS server configuration details for a specific SVM

```
# The API:
GET
/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_i
d}

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-
005056ae78de/10878444899913433090/103985" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>"

# The response:
{
"node": {
  "uuid": "a5f65ec0-3550-11ee-93c5-005056ae78de",
  "name": "sti220-vsim-sr050u"
},
"svm": {
  "uuid": "80e795f4-3553-11ee-9f97-005056ae78de",
  "name": "vs0"
},
"identifier": 10878444899913433000,
"connection_id": 103985,
"server_ip": "172.29.212.166",
"client_ip": "10.249.30.206",
"authentication": "ntlmv2",
"user": "example_user",
"mapped_unix_user": "example_user",
"open_shares": 1,
"open_files": 5,
"open_other": 0,
"connected_duration": "PT50M20S",
"idle_duration": "PT5M41S",
"protocol": "smb3_1",
"continuous_availability": "unavailable",
"smb_signing": true,
"smb_encryption": "unencrypted",
```

```
"connection_count": 1,
"large_mtu": true,
"volumes": [
  {
    "name": "root_vs0",
    "uuid": "5015c99b-3554-11ee-9f97-005056ae78de"
  },
  {
    "name": "vol_test",
    "uuid": "84b94bb8-3553-11ee-a3c3-005056ae0dd5"
  }
]
}
```

#### Removing all existing CIFS sessions for a specific node on a specific SVM

To delete all the existing CIFS session, pass the identifier and connection ID as zero (0) in the following API. This will delete all of the CIFS sessions for the given SVM on the node.

```
# The API:
DELETE
/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_i
d}

# The call:
curl -X DELETE "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-005056ae78de/0/0" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>"
```

#### Removing all CIFS sessions for a specific connection on a specific node on a specific SVM

To delete a CIFS session, use the following API. This will delete the CIFS sessions for a given SVM on the node.

```
# The API:
DELETE
/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_i
d}

# The call:
curl -X DELETE "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-005056ae78de/0/103985" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>"
```

---

#### Removing a specific CIFS session for a specific Node on a specific SVM

To delete a specific CIFS session, use the following API. This will delete the specific CIFS session for the given SVM on the node.

---

```
# The API:
DELETE
/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_i
d}

# The call:
curl -X DELETE "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions/a5f65ec0-3550-11ee-93c5-
005056ae78de/80e795f4-3553-11ee-9f97-
005056ae78de/10878444899913433090/103985" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>"
```

---

#### Retrieve CIFS session information for all SVMs

GET `/protocols/cifs/sessions`

**Introduced In:** 9.8

Retrieves the CIFS sessions information for all SVMs.

##### Related ONTAP commands

- `vserver cifs session show -active-volumes`

**Learn more**

- DOC /protocols/cifs/sessions

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| smb_signing | boolean | query | False | Filter by smb_signing |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| mapped_unix_user | string | query | False | Filter by mapped_unix_user |
| open_shares | integer | query | False | Filter by open_shares |
| continuous_availability | string | query | False | Filter by continuous_availability |
| open_other | integer | query | False | Filter by open_other |
| protocol | string | query | False | Filter by protocol |
| open_files | integer | query | False | Filter by open_files |
| user | string | query | False | Filter by user |
| identifier | integer | query | False | Filter by identifier |
| large_mtu | boolean | query | False | Filter by large_mtu |
| server_ip | string | query | False | Filter by server_ip |
| connected_duration | string | query | False | Filter by connected_duration |
| authentication | string | query | False | Filter by authentication |
| idle_duration | string | query | False | Filter by idle_duration |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| smb_encryption | string | query | False | Filter by smb_encryption |
| client_ip | string | query | False | Filter by client_ip |
| connection_id | integer | query | False | Filter by connection_id |
| connection_count | integer | query | False | Filter by connection_count |
| volumes.name | string | query | False | Filter by volumes.name |
| volumes.uuid | string | query | False | Filter by volumes.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_session] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication": "ntlmv2",
      "client_ip": "10.74.7.182",
      "connected_duration": "P4DT84H30M5S",
      "connection_count": 0,
      "connection_id": 22802,
      "continuous_availability": "unavailable",
      "identifier": 4622663542519103507,
      "idle_duration": "P4DT84H30M5S",
      "large_mtu": true,
      "mapped_unix_user": "root",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "open_files": 0,
      "open_other": 0,
      "open_shares": 0,
      "protocol": "smb3_1",
      "server_ip": "10.140.78.248",
      "smb_encryption": "unencrypted",
      "smb_signing": "",
      "svm": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "user": "NBCIFSQA2\\administrator",
      "volumes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ]
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volumes

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.8<br>• x-nullable: true |

cifs_session

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication | string | SMB authentication over which the client accesses the share. The following values are supported:<br>• none - No authentication<br>• ntlmv1 - Ntlm version 1 mechanism<br>• ntlmv2 - Ntlm version 2 mechanism<br>• kerberos - Kerberos authentication<br>• anonymous - Anonymous mechanism |
| client_ip | string | Specifies IP address of the client. |
| connected_duration | string | Specifies an ISO-8601 format of date and time used to retrieve the connected time duration in hours, minutes, and seconds format. |
| connection_count | integer | A counter used to track requests that are sent to the volumes to the node. |

| Name | Type | Description |
| --- | --- | --- |
| connection_id | integer | A unique 32-bit unsigned number used to represent each SMB session's connection ID. |
| continuous_availability | string | The level of continuous availability protection provided to the SMB sessions and/or files.<br>• unavailable - Open file is not continuously available. For sessions, it contains one or more open files but none of them are continuously available.<br>• available - open file is continuously available. For sessions, it contains one or more open files and all of them are continuously available.<br>• partial - Sessions only. Contains at least one continuously available open file with other files open but not continuously available. |
| identifier | integer | A unique 64-bit unsigned number used to represent each SMB session's identifier. |
| idle_duration | string | Specifies an ISO-8601 format of date and time used to retrieve the idle time duration in hours, minutes, and seconds format. |
| large_mtu | boolean | Specifies whether the large MTU is enabled or not for an SMB session. |
| mapped_unix_user | string | Indicated that a mapped UNIX user has logged in. |
| node | node | |
| open_files | integer | Number of files the SMB session has opened. |

| Name | Type | Description |
| --- | --- | --- |
| open_other | integer | Number of other files the SMB session has opened. |
| open_shares | integer | Number of shares the SMB session has opened. |
| protocol | string | The SMB protocol version over which the client accesses the volumes. The following values are supported:<br>• smb1 - SMB version 1<br>• smb2 - SMB version 2<br>• smb2_1 - SMB version 2 minor version 1<br>• smb3 - SMB version 3<br>• smb3_1 - SMB version 3 minor version 1 |
| server_ip | string | Specifies the IP address of the SVM. |
| smb_encryption | string | Indicates an SMB encryption state. The following values are supported:<br>• unencrypted - SMB session is not encrypted<br>• encrypted - SMB session is fully encrypted. SVM level encryption is enabled and encryption occurs for the entire session.<br>• partially_encrypted - SMB session is partially encrypted. Share level encryption is enabled and encryption is initiated when the tree-connect occurs. |
| smb_signing | boolean | Specifies whether or not SMB signing is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
|---|---|---|
| user | string | Indicates that a Windows user has logged in. |
| volumes | array[volumes] | A group of volumes, the client is accessing. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete SMB session information for an SVM node**

DELETE
`/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_id}`

**Introduced In:** 9.9

Deletes SMB session information on a node for an SVM.

- To delete the specific SMB session information, pass the relevant SMB session's identifier and connection Id.
- To delete all the SMB session information on specific node and SVM, pass the both SMB session's identifier and connection Id as zero(0)
- To delete all the SMB session information on specific connection, pass the specific SMB session's Identifier value as zero(0).
- To delete all the SMB session information on specific Identifier alone is not allowed.

#### Related ONTAP commands

- `vserver cifs session close`

#### Learn more

- DOC /protocols/cifs/sessions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID. |
| identifier | integer | path | True | Unique identifier for the SMB session. |
| connection_id | integer | path | True | Unique identifier for the SMB connection. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655503 | The SMB session delete does not allow a connection ID of zero (0). |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve SMB session information for an SMB connection in an SVM node

GET `/protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_id}`

**Introduced In:** 9.8

Retrieves specific SMB session information for a specific SMB connection in a node on an SVM.

#### Learn more

- DOC /protocols/cifs/sessions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID. |
| identifier | integer | path | True | Unique identifier for the SMB session. |
| connection_id | integer | path | True | Unique identifier for the SMB connection. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication | string | SMB authentication over which the client accesses the share. The following values are supported:<br>• none - No authentication<br>• ntlmv1 - Ntlm version 1 mechanism<br>• ntlmv2 - Ntlm version 2 mechanism<br>• kerberos - Kerberos authentication<br>• anonymous - Anonymous mechanism |

| Name | Type | Description |
| --- | --- | --- |
| client_ip | string | Specifies IP address of the client. |
| connected_duration | string | Specifies an ISO-8601 format of date and time used to retrieve the connected time duration in hours, minutes, and seconds format. |
| connection_count | integer | A counter used to track requests that are sent to the volumes to the node. |
| connection_id | integer | A unique 32-bit unsigned number used to represent each SMB session's connection ID. |
| continuous_availability | string | The level of continuous availability protection provided to the SMB sessions and/or files.<br>• unavailable - Open file is not continuously available. For sessions, it contains one or more open files but none of them are continuously available.<br>• available - open file is continuously available. For sessions, it contains one or more open files and all of them are continuously available.<br>• partial - Sessions only. Contains at least one continuously available open file with other files open but not continuously available. |
| identifier | integer | A unique 64-bit unsigned number used to represent each SMB session's identifier. |
| idle_duration | string | Specifies an ISO-8601 format of date and time used to retrieve the idle time duration in hours, minutes, and seconds format. |
| large_mtu | boolean | Specifies whether the large MTU is enabled or not for an SMB session. |

| Name | Type | Description |
| --- | --- | --- |
| mapped_unix_user | string | Indicated that a mapped UNIX user has logged in. |
| node | node | |
| open_files | integer | Number of files the SMB session has opened. |
| open_other | integer | Number of other files the SMB session has opened. |
| open_shares | integer | Number of shares the SMB session has opened. |
| protocol | string | The SMB protocol version over which the client accesses the volumes. The following values are supported:<br>• smb1 - SMB version 1<br>• smb2 - SMB version 2<br>• smb2_1 - SMB version 2 minor version 1<br>• smb3 - SMB version 3<br>• smb3_1 - SMB version 3 minor version 1 |
| server_ip | string | Specifies the IP address of the SVM. |
| smb_encryption | string | Indicates an SMB encryption state. The following values are supported:<br>• unencrypted - SMB session is not encrypted<br>• encrypted - SMB session is fully encrypted. SVM level encryption is enabled and encryption occurs for the entire session.<br>• partially_encrypted - SMB session is partially encrypted. Share level encryption is enabled and encryption is initiated when the tree-connect occurs. |

| Name | Type | Description |
| --- | --- | --- |
| smb_signing | boolean | Specifies whether or not SMB signing is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| user | string | Indicates that a Windows user has logged in. |
| volumes | array[volumes] | A group of volumes, the client is accessing. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": "ntlmv2",
  "client_ip": "10.74.7.182",
  "connected_duration": "P4DT84H30M5S",
  "connection_count": 0,
  "connection_id": 22802,
  "continuous_availability": "unavailable",
  "identifier": 4622663542519103507,
  "idle_duration": "P4DT84H30M5S",
  "large_mtu": true,
  "mapped_unix_user": "root",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "open_files": 0,
  "open_other": 0,
  "open_shares": 0,
  "protocol": "smb3_1",
  "server_ip": "10.140.78.248",
  "smb_encryption": "unencrypted",
  "smb_signing": "",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user": "NBCIFSQA2\\administrator",
  "volumes": [
    {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volumes

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.8<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update shadow copy information

### View and update shadow copy information

Microsoft Remote Volume Shadow Copy Services (VSS) is an extension of the existing Microsoft VSS infrastructure. Previously, VSS could be used for backup services only for data stored on the local disk. This limited the use of VSS to applications that stored data either on a local disk or on SAN-based storage. With Remote VSS, Microsoft has extended the VSS infrastructure to support the shadow copying of SMB shares. Server applications, such as, Hyper-V are now storing VHD files on SMB file shares. With these new extensions, it is possible to take application consistent shadow copies for virtual machines that store data and configuration files on shares.

#### Retrieving Shadow copy sets for all SVMs

```
# The API:
GET /api/protocols/cifs/shadowcopy-sets

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shadowcopy-
sets?fields=*&return_records=true&return_timeout=15" -H  "accept:
application/json" -H  "authorization: Basic <CREDENTIALS>"

# The Response:
{
"records": [
  {
    "uuid": "27ed1b79-97f6-11ec-8ad1-0050568e57be",
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "keep_snapshots": true
  },
  {
    "uuid": "388be551-97f6-11ec-8ad1-0050568e57be",
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "keep_snapshots": false
  },
  {
    "uuid": "525104ef-9f96-11ec-82fd-0050568e57be",
    "svm": {
      "uuid": "fdb5bd8b-9498-11ec-9f9d-0050568e57be",
      "name": "vs2"
    },
    "keep_snapshots": false
  },
  {
    "uuid": "66f8f723-9f96-11ec-82fd-0050568e57be",
    "svm": {
      "uuid": "fdb5bd8b-9498-11ec-9f9d-0050568e57be",
      "name": "vs2"
    },
    "keep_snapshots": true
```

```
  }
],
"num_records": 4
}
```

#### Retrieving information for a specific shadow copy set

```
# The API:
GET /api/protocols/cifs/shadowcopy-sets/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shadowcopy-
sets/525104ef-9f96-11ec-82fd-0050568e57be?fields=*" -H  "accept:
application/json"

# The Response:
{
"uuid": "525104ef-9f96-11ec-82fd-0050568e57be",
"svm": {
  "uuid": "fdb5bd8b-9498-11ec-9f9d-0050568e57be",
  "name": "vs2"
},
"keep_snapshots": false
}
```

#### Updating the keep-snapshot property of a specific shadow copy set

```
# The API:
GET /api/protocols/cifs/shadowcopy-sets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shadowcopy-
sets/525104ef-9f96-11ec-82fd-0050568e57be" -H  "accept: application/json"
-H  "Content-Type: application/json" -d "{  \"keep_snapshots\": true}"
```

#### Retrieving shadow copy information for all SVMs

```
# The API:
GET /api/protocols/cifs/shadow-copies

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shadow-
copies?fields=*&return_records=true&return_timeout=15" -H  "accept:
application/json"

# The Response:
{
"records": [
  {
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "shadowcopy_set": {
      "uuid": "9169cd4a-a421-11ec-b82e-0050568e57be"
    },
    "uuid": "919fbc80-a421-11ec-b82e-0050568e57be",
    "share": {
      "name": "sh1"
    },
    "client_uuid": "918536b7-a421-11ec-b82e-0050568e57be",
    "volume": {
      "uuid": "e117c8f6-9498-11ec-9f9d-0050568e57be",
      "name": "vol1"
    }
  },
  {
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "shadowcopy_set": {
      "uuid": "9169cd4a-a421-11ec-b82e-0050568e57be"
    },
    "uuid": "91ac5a5f-a421-11ec-b82e-0050568e57be",
    "share": {
      "name": "sh2"
    },
    "client_uuid": "91a9252b-a421-11ec-b82e-0050568e57be",
```

```
    "volume": {
      "uuid": "e117c8f6-9498-11ec-9f9d-0050568e57be",
      "name": "vol1"
    }
  },
  {
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "shadowcopy_set": {
      "uuid": "9169cd4a-a421-11ec-b82e-0050568e57be"
    },
    "uuid": "91b14098-a421-11ec-b82e-0050568e57be",
    "share": {
      "name": "sh3"
    },
    "client_uuid": "91adfc85-a421-11ec-b82e-0050568e57be",
    "volume": {
      "uuid": "e117c8f6-9498-11ec-9f9d-0050568e57be",
      "name": "vol1"
    }
  },
  {
    "svm": {
      "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
      "name": "vs1"
    },
    "shadowcopy_set": {
      "uuid": "9169cd4a-a421-11ec-b82e-0050568e57be"
    },
    "uuid": "91b63309-a421-11ec-b82e-0050568e57be",
    "share": {
      "name": "sh4"
    },
    "client_uuid": "91b2f817-a421-11ec-b82e-0050568e57be",
    "volume": {
      "uuid": "e117c8f6-9498-11ec-9f9d-0050568e57be",
      "name": "vol1"
    }
  }
],
"num_records": 4
}
```

### Retrieving information for a specific shadow copy

```
# The API:
GET /api/protocols/cifs/shadow-copies/{client_uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shadow-copies/91adfc85-
a421-11ec-b82e-0050568e57be?fields=*" -H  "accept: application/json"

# The Response:
{
"svm": {
  "uuid": "dfb8e00d-9498-11ec-9f9d-0050568e57be",
  "name": "vs1"
},
"shadowcopy_set": {
  "uuid": "9169cd4a-a421-11ec-b82e-0050568e57be"
},
"uuid": "91b14098-a421-11ec-b82e-0050568e57be",
"share": {
  "name": "sh3"
},
"client_uuid": "91adfc85-a421-11ec-b82e-0050568e57be",
"volume": {
  "uuid": "e117c8f6-9498-11ec-9f9d-0050568e57be",
  "name": "vol1"
}
}
```

### Updating the list of files to be shadowcopied in a particular share

Use this endpoint to update the list of files to be shadow copied in a particular share. Set "restore" field as false to perform this operation.

```
# The API:
PATCH /protocols/cifs/shadow-copies/{client_uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shadow-
copies/91adfc85-a421-11ec-b82e-0050568e57be?restore=false" -H  "accept:
application/json" -d "{  \"files\": [    \"/vhd1\",    \"/vhd2\"  ]}"
```

---

#### Requesting the storage system to restore a directory

You set restore field to true to perform this operation. Only users with the security login role "vsadmin" can perform the operation.

---

```
# The API:
PATCH /protocols/cifs/shadow-copies?restore=true

# The call:
PATCH "api/protocols/cifs/shadow-copies?restore=true" -d "{
\"destination_dir\": \"/dir2\",  \"source_dir\": \"/src_dir\",
\"volume\": {    \"name\": \"test_vol\"  },  \"with_content\": false}"
```

---

### Retrieve shadow copies

GET `/protocols/cifs/shadow-copies`

**Introduced In:** 9.11

Retrieves Shadowcopies

#### Related ONTAP commands

- `vserver cifs shadowcopy show-shares`

#### Learn more

- DOC /protocols/cifs/shadow-copies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share.name | string | query | False | Filter by share.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| files | string | query | False | Filter by files |
| client_uuid | string | query | False | Filter by client_uuid |
| uuid | string | query | False | Filter by uuid |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| shadowcopy_set.uuid | string | query | False | Filter by shadowcopy_set.uuid |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |
| Limit the number of records returned. | return_records | boolean | query | False |
| The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 | return_timeout | integer | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |

| Name | Type | Description |
| --- | --- | --- |
| records | array[shadowcopy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_uuid": "abc13450-1f19-40ba-9b82-ebf277517e7e",
      "destination_dir": "/dir2",
      "files": [
        "/vhd1",
        "/vhd2"
      ],
      "shadowcopy_set": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "5f4e19f3-5e2d-11ec-8766-005056a7c134"
      },
      "share": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "sh1"
      },
      "source_dir": "/dir1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "fef32805-1f19-40ba-9b82-ebf277517e7e",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

shadowcopy_set

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| uuid | string | |

share

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Share name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

The name of the volume where the source and destination directories reside.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

shadowcopy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| client_uuid | string | The universally-unique identifier of the client's shadow copy. |
| destination_dir | string | The path of the destination directory. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| files | array[string] | The list of files to shadow copy in the share. The path is in UTF8 and uses forward slash as a directory separator. The path is relative to the root of the share. |
| shadowcopy_set | shadowcopy_set | |
| share | share | |
| source_dir | string | The path of the source directory. The path is in UTF8 and uses forward slash as a directory separator. The path is relative to the root of the volume. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy. |
| volume | volume | The name of the volume where the source and destination directories reside. |
| with_content | boolean | Specifies what needs to be restored. False specifies the directory only. True indicates the directory and its content. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a shadow copy

GET `/protocols/cifs/shadow-copies/{client_uuid}`

**Introduced In:** 9.11

Retrieves a Shadowcopy

#### Related ONTAP commands

- `vserver cifs shadowcopy show-shares`

#### Learn more

- DOC /protocols/cifs/shadow-copies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| client_uuid | string | path | True | client shadowcopy ID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| client_uuid | string | The universally-unique identifier of the client's shadow copy. |

| Name | Type | Description |
| --- | --- | --- |
| destination_dir | string | The path of the destination directory. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the volume. |
| files | array[string] | The list of files to shadow copy in the share. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the share. |
| shadowcopy_set | shadowcopy_set | |
| share | share | |
| source_dir | string | The path of the source directory. The path is in UTF8 and uses forward slash<br>as a directory separator. The path is relative to the root of the volume. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy. |
| volume | volume | The name of the volume where the source and destination directories reside. |
| with_content | boolean | Specifies what needs to be restored. False specifies the directory only.<br>True indicates the directory and its content. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "client_uuid": "abc13450-1f19-40ba-9b82-ebf277517e7e",
  "destination_dir": "/dir2",
  "files": [
    "/vhd1",
    "/vhd2"
  ],
  "shadowcopy_set": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "5f4e19f3-5e2d-11ec-8766-005056a7c134"
  },
  "share": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "sh1"
  },
  "source_dir": "/dir1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "fef32805-1f19-40ba-9b82-ebf277517e7e",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
```

```
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

shadowcopy_set

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| uuid | string | |

share

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Share name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

The name of the volume where the source and destination directories reside.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update a shadow copy**

PATCH `/protocols/cifs/shadow-copies/{client_uuid}`

**Introduced In:** 9.11

Specify list Files to be added as a part of Shadowcopy creation

**Learn more**

- DOC /protocols/cifs/shadow-copies

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| client_uuid | string | path | True | client shadowcopy ID |
| restore | boolean | query | False | Indicates whether to restore a directory |

**Request Body**

| Name | Type | Description |
|---|---|---|
| client_uuid | string | The universally-unique identifier of the client's shadow copy. |
| destination_dir | string | The path of the destination directory. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the volume. |
| files | array[string] | The list of files to shadow copy in the share. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the share. |
| shadowcopy_set | shadowcopy_set | |
| share | share | |
| source_dir | string | The path of the source directory. The path is in UTF8 and uses forward slash<br>as a directory separator. The path is relative to the root of the volume. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The universally-unique identifier of the storage's shadow copy. |
| volume | volume | The name of the volume where the source and destination directories reside. |
| with_content | boolean | Specifies what needs to be restored. False specifies the directory only.<br>True indicates the directory and its content. |

**Example request**

```
{
  "client_uuid": "abc13450-1f19-40ba-9b82-ebf277517e7e",
  "destination_dir": "/dir2",
  "files": [
    "/vhd1",
    "/vhd2"
  ],
  "shadowcopy_set": {
    "uuid": "5f4e19f3-5e2d-11ec-8766-005056a7c134"
  },
  "share": {
    "name": "sh1"
  },
  "source_dir": "/dir1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "fef32805-1f19-40ba-9b82-ebf277517e7e",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| storage_shadowcopy_set_uuid | string | The universally-unique identifier of the storage's shadow copy set. |
| storage_shadowcopy_uuid | string | The universally-unique identifier of the storage's shadow copy. |

**Example response**

```
{
  "storage_shadowcopy_set_uuid": "f8328660-00e6-11e6-80d9-
005056bd65a9",
  "storage_shadowcopy_uuid": "fef32805-1f19-40ba-9b82-ebf277517e7e"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262196 | The 'svm' field is not supported in the body of a PATCH request. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

self_link

shadowcopy_set

| Name | Type | Description |
|---|---|---|
| uuid | string | |

share

| Name | Type | Description |
|---|---|---|
| name | string | Share name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

The name of the volume where the source and destination directories reside.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

shadowcopy

| Name | Type | Description |
| --- | --- | --- |
| client_uuid | string | The universally-unique identifier of the client's shadow copy. |
| destination_dir | string | The path of the destination directory. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the volume. |
| files | array[string] | The list of files to shadow copy in the share. The path is in UTF8 and uses forward<br>slash as a directory separator. The path is relative to the root of the share. |
| shadowcopy_set | shadowcopy_set | |
| share | share | |
| source_dir | string | The path of the source directory. The path is in UTF8 and uses forward slash<br>as a directory separator. The path is relative to the root of the volume. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy. |

| Name | Type | Description |
|---|---|---|
| volume | volume | The name of the volume where the source and destination directories reside. |
| with_content | boolean | Specifies what needs to be restored. False specifies the directory only.<br>True indicates the directory and its content. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve shadow copy sets

GET `/protocols/cifs/shadowcopy-sets`

**Introduced In:** 9.11

Retrieves Shadowcopy Sets.

#### Related ONTAP commands

- `vserver cifs shadowcopy show-sets`

#### Learn more

- DOC /protocols/cifs/shadow-copies

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | query | False | Filter by uuid |
| keep_snapshots | boolean | query | False | Filter by keep_snapshots |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |
| Limit the number of records returned. | return_records | boolean | query | False |
| The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 | return_timeout | integer | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[shadowcopy_set] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "keep_snapshots": 1,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "f8328660-00e6-11e6-80d9-005056bd65a9"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

shadowcopy_set

| Name | Type | Description |
|---|---|---|
| keep_snapshots | boolean | Request the storage system to keep the snapshot copies taken as a part of the shadow copy set creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy set. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a shadow copy set

GET `/protocols/cifs/shadowcopy-sets/{uuid}`

**Introduced In:** 9.11

Retrieves a Shadowcopy set

#### Related ONTAP commands

- `vserver cifs shadowcopy show-sets`

#### Learn more

- DOC /protocols/cifs/shadow-copies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Storage shadowcopy set ID. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| keep_snapshots | boolean | Request the storage system to keep the snapshot copies taken as a part of the shadow copy set creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy set. |

#### Example response

```
{
  "keep_snapshots": 1,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "f8328660-00e6-11e6-80d9-005056bd65a9"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a shadow copy set

PATCH `/protocols/cifs/shadowcopy-sets/{uuid}`

**Introduced In:** 9.11

Updates a Shadowcopy set

#### Learn more

- DOC /protocols/cifs/shadowcopy

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Storage shadowcopy set ID. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| keep_snapshots | boolean | Request the storage system to keep the snapshot copies taken as a part of the shadow copy set creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy set. |

#### Example request

```
{
  "keep_snapshots": 1,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "f8328660-00e6-11e6-80d9-005056bd65a9"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262196 | The 'svm' field is not supported in the body of a PATCH request. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

shadowcopy_set

| Name | Type | Description |
|---|---|---|
| keep_snapshots | boolean | Request the storage system to keep the snapshot copies taken as a part of the shadow copy set creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The universally-unique identifier of the storage's shadow copy set. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage CIFS shares

### Manage CIFS shares

#### Overview

Before any users or applications can access data on the CIFS server over SMB, a CIFS share must be created with sufficient share permissions. CIFS share is a named access point in a volume which is tied to the CIFS server on the SVM. Before creating a CIFS share make sure that the path is valid within the scope of the SVM and that it is reachable.

Permissions can be assigned to this newly created share by specifying the 'acls' field. When a CIFS share is created, ONTAP creates a default ACL for this share with 'Full-Control' permissions for an 'Everyone' user.

#### Examples

### Creating a CIFS share

To create a CIFS share for a CIFS server, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/shares

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/shares?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"access_based_enumeration\": false, \"acls\": [ { \"permission\":
\"no_access\", \"type\": \"unix_user\", \"user_or_group\": \"root\" } ],
\"change_notify\": true, \"comment\": \"HR Department Share\",
\"encryption\": false, \"home_directory\": false, \"name\": \"TEST\",
\"oplocks\": true, \"no_strict_security\": true, \"path\": \"/\", \"svm\":
{ \"name\": \"vs1\", \"uuid\": \"6d8e8870-8753-11eb-8d86-0050568ea61a\" },
\"unix_symlink\": \"local\", \"show_snapshot\": \"true\",
```

```
\"continuously_available\": \"false\", \"namespace_caching\": \"true\",
\"allow_unencrypted_access\": \"false\", \"file_umask\": \"025\",
\"dir_umask\": \"026\", \"offline_files\": \"documents\",
\"vscan_profile\": \"standard\", \"force_group_for_create\": \"root\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
      "name": "vs1"
    },
    "name": "TEST",
    "path": "/",
    "comment": "HR Department Share",
    "home_directory": false,
    "oplocks": true,
    "access_based_enumeration": false,
    "change_notify": true,
    "encryption": false,
    "unix_symlink": "local",
    "acls": [
      {
        "user_or_group": "root",
        "type": "unix_user",
        "permission": "no_access",
        "win_sid_unix_id": "0"
      }
    ],
    "no_strict_security": true,
    "show_snapshot": true,
    "continuously_available": false,
    "namespace_caching": true,
    "allow_unencrypted_access": false,
    "file_umask": "25",
    "dir_umask": "26",
    "offline_files": "documents",
    "vscan_profile": "standard",
    "force_group_for_create": "root"
  }
]
}
```

**Retrieving CIFS Shares for all SVMs in the cluster**

```
# The API:
GET /api/protocols/cifs/shares

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/shares?fields=*&return_records=true&return_timeout=
15" -H "accept application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6d8e8870-8753-11eb-8d86-0050568ea61a"
        }
      }
    },
    "name": "c$",
    "path": "/",
    "home_directory": false,
    "oplocks": true,
    "access_based_enumeration": false,
    "change_notify": true,
    "encryption": false,
    "unix_symlink": "local",
    "acls": [
      {
        "user_or_group": "BUILTIN\\Administrators",
        "type": "windows",
        "permission": "full_control",
        "win_sid_unix_id": "S-1-1-0"
      }
    ],
    "volume": {
      "name": "vol1",
      "uuid": "6f4fb33a-8753-11eb-8d86-0050568ea61a"
    },
    "show_snapshot": false,
```

```
    "cache_attribute": false,
    "continuously_available": false,
    "namespace_caching": false,
    "allow_unencrypted_access": true,
    "vscan_profile": "standard",
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/shares/6d8e8870-8753-11eb-8d86-
0050568ea61a/c%24"
      }
    }
  },
  {
    "svm": {
      "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6d8e8870-8753-11eb-8d86-0050568ea61a"
        }
      }
    },
    "name": "ipc$",
    "path": "/",
    "home_directory": false,
    "oplocks": false,
    "access_based_enumeration": false,
    "change_notify": false,
    "encryption": false,
    "volume": {
      "name": "vol1",
      "uuid": "6f4fb33a-8753-11eb-8d86-0050568ea61a"
    },
    "show_snapshot": false,
    "cache_attribute": false,
    "continuously_available": false,
    "namespace_caching": false,
    "allow_unencrypted_access": false,
    "vscan_profile": "standard",
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/shares/6d8e8870-8753-11eb-8d86-
0050568ea61a/ipc%24"
      }
    }
  },
```

```
  {
    "svm": {
      "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6d8e8870-8753-11eb-8d86-0050568ea61a"
        }
      }
    },
    "name": "TEST",
    "path": "/",
    "comment": "HR Department Share",
    "home_directory": false,
    "oplocks": true,
    "access_based_enumeration": false,
    "change_notify": true,
    "encryption": false,
    "unix_symlink": "local",
    "acls": [
      {
        "user_or_group": "Everyone",
        "type": "windows",
        "permission": "full_control",
        "win_sid_unix_id": "S-1-1-0"
      },
      {
        "user_or_group": "root",
        "type": "unix_user",
        "permission": "no_access",
        "win_sid_unix_id": "0"
      }
    ],
    "volume": {
      "name": "vol1",
      "uuid": "6f4fb33a-8753-11eb-8d86-0050568ea61a"
    },
    "no_strict_security": true,
    "show_snapshot": true,
    "cache_attribute": true,
    "continuously_available": false,
    "namespace_caching": true,
    "allow_unencrypted_access": true,
    "file_umask": "25",
    "dir_umask": "26",
    "offline_files": "documents",
```

```
    "vscan_profile": "standard",
    "force_group_for_create": "root",
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/shares/6d8e8870-8753-11eb-8d86-
0050568ea61a/TEST"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href":
"/api/protocols/cifs/shares?fields=*&return_records=true&return_timeout=15
"
  }
}
}
```

#### Retrieving all CIFS Shares for all SVMs in the cluster for which the acls are configured for a "root" user

```
# The API:
GET /api/protocols/cifs/shares

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/shares?acls.user_or_group=root&fields=*&return_reco
rds=true&return_timeout=15" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
      "name": "vs1"
    },
    "name": "TEST",
    "path": "/",
    "comment": "HR Department Share",
    "home_directory": false,
```

```
    "oplocks": true,
    "access_based_enumeration": false,
    "change_notify": true,
    "encryption": false,
    "unix_symlink": "local",
    "acls": [
      {
        "user_or_group": "Everyone",
        "type": "windows",
        "permission": "full_control",
        "win_sid_unix_id": "S-1-1-0"
      },
      {
        "user_or_group": "root",
        "type": "unix_user",
        "permission": "no_access",
        "win_sid_unix_id": "0"
      }
    ],
    "volume": {
      "name": "vol1",
      "uuid": "6f4fb33a-8753-11eb-8d86-0050568ea61a"
    },
    "no_strict_security": true,
    "show_snapshot": true,
    "cache_attribute": true,
    "continuously_available": false,
    "namespace_caching": true,
    "allow_unencrypted_access": true,
    "file_umask": "25",
    "dir_umask": "26",
    "offline_files": "documents",
    "vscan_profile": "standard",
    "force_group_for_create": "root"
  }
],
"num_records": 1
}
```


**Retrieving a specific CIFS share configuration for an SVM**

The configuration being returned is identified by the UUID of its SVM and the name of the share.

```
# The API:
GET /api/protocols/cifs/shares/{svm.uuid}/{name}
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/6d8e8870-8753-
11eb-8d86-0050568ea61a/TEST" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "6d8e8870-8753-11eb-8d86-0050568ea61a",
  "name": "vs1"
},
"name": "TEST",
"path": "/",
"comment": "HR Department Share",
"home_directory": false,
"oplocks": true,
"access_based_enumeration": false,
"change_notify": true,
"encryption": false,
"unix_symlink": "local",
"acls": [
  {
    "user_or_group": "Everyone",
    "type": "windows",
    "permission": "full_control",
    "win_sid_unix_id": "S-1-1-0"
  },
  {
    "user_or_group": "root",
    "type": "unix_user",
    "permission": "no_access",
    "win_sid_unix_id": "0"
  }
],
"volume": {
  "name": "vol1",
  "uuid": "6f4fb33a-8753-11eb-8d86-0050568ea61a"
},
"no_strict_security": true,
"show_snapshot": true,
"cache_attribute": true,
"continuously_available": false,
"namespace_caching": true,
"allow_unencrypted_access": true,
"file_umask": "25",
"dir_umask": "26",
"offline_files": "documents",
```

```
"vscan_profile": "standard",
"force_group_for_create": "root"
}
```

#### Updating a specific CIFS share for an SVM

The CIFS share being modified is identified by the UUID of its SVM and the CIFS share name. The CIFS share ACLs cannot be modified with this API.

```
# The API:
PATCH /api/protocols/cifs/shares/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shares/6d8e8870-8753-
11eb-8d86-0050568ea61a/TEST" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"access_based_enumeration\": true,
\"change_notify\": true, \"comment\": \"HR Department Share\",
\"encryption\": false, \"oplocks\": true, \"no_strict_security\": true,
\"path\": \"/\", \"unix_symlink\": \"widelink\", \"show_snapshot\":
\"false\", \"continuously_available\": \"true\", \"namespace_caching\":
\"false\", \"allow_unencrypted_access\": \"true\", \"file_umask\":
\"022\", \"dir_umask\": \"022\", \"offline_files\": \"programs\",
\"vscan_profile\": \"no_scan\", \"force_group_for_create\": \"root\"}"
```

#### Removing a specific CIFS share for an SVM

The CIFS share being removed is identified by the UUID of its SVM and the CIFS share name.

```
# The API:
DELETE /api/protocols/cifs/shares/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/shares/6d8e8870-8753-
11eb-8d86-0050568ea61a/test" -H "accept: application/json"
```

---

### Retrieve CIFS shares

GET `/protocols/cifs/shares`

**Introduced In:** 9.6

Retrieves CIFS shares.

#### Related ONTAP commands

- `vserver cifs share show`
- `vserver cifs share properties show`

#### Learn more

- DOC /protocols/cifs/shares

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| offline_files | string | query | False | Filter by offline_files<br>• Introduced in: 9.10 |
| browsable | boolean | query | False | Filter by browsable<br>• Introduced in: 9.13 |
| allow_unencrypted_access | boolean | query | False | Filter by allow_unencrypted_access<br>• Introduced in: 9.11 |
| encryption | boolean | query | False | Filter by encryption |
| max_connections_per_share | integer | query | False | Filter by max_connections_per_share<br>• Introduced in: 9.14 |
| home_directory | boolean | query | False | Filter by home_directory |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| namespace_caching | boolean | query | False | Filter by namespace_caching<br><br>• Introduced in: 9.10 |
| acls.type | string | query | False | Filter by acls.type |
| acls.user_or_group | string | query | False | Filter by acls.user_or_group |
| acls.permission | string | query | False | Filter by acls.permission |
| acls.win_sid_unix_id | string | query | False | Filter by acls.win_sid_unix_id<br><br>• Introduced in: 9.16 |
| vscan_profile | string | query | False | Filter by vscan_profile<br><br>• Introduced in: 9.10 |
| continuously_available | boolean | query | False | Filter by continuously_available<br><br>• Introduced in: 9.10 |
| show_snapshot | boolean | query | False | Filter by show_snapshot<br><br>• Introduced in: 9.10 |
| oplocks | boolean | query | False | Filter by oplocks |
| access_based_enumeration | boolean | query | False | Filter by access_based_enumeration |
| change_notify | boolean | query | False | Filter by change_notify |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• maxLength: 80<br>• minLength: 1 |
| show_previous_versions | boolean | query | False | Filter by show_previous_versions<br>• Introduced in: 9.13 |
| dir_umask | string | query | False | Filter by dir_umask<br>• Introduced in: 9.10 |
| attribute_cache | boolean | query | False | Filter by attribute_cache<br>• Introduced in: 9.14 |
| path | string | query | False | Filter by path<br>• maxLength: 256<br>• minLength: 1 |
| unix_symlink | string | query | False | Filter by unix_symlink |
| file_umask | string | query | False | Filter by file_umask<br>• Introduced in: 9.10 |
| no_strict_security | boolean | query | False | Filter by no_strict_security<br>• Introduced in: 9.9 |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 1 |
| force_group_for_create | string | query | False | Filter by force_group_for_create<br>• Introduced in: 9.10 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_share] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "acls": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "permission": "string",
          "type": "string",
          "user_or_group": "ENGDOMAIN\\ad_user",
          "win_sid_unix_id": "string"
        }
      ],
      "comment": "HR Department Share",
      "dir_umask": "21",
      "file_umask": "21",
      "force_group_for_create": "string",
      "name": "HR_SHARE",
      "offline_files": "string",
      "path": "/volume_1/eng_vol/",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
```

```
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_symlink": "string",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "vscan_profile": "string"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |
| win_sid_unix_id | string | Windows SID/UNIX ID depending on access-control type. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |

| Name | Type | Description |
| --- | --- | --- |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value:<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group (also called the "force-group").<br>The "force-group" must be a predefined group in the UNIX group<br>database. This setting has no effect unless the security style of the volume is UNIX or mixed security style. |

| Name | Type | Description |
| --- | --- | --- |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an<br>individual SMB share for each user.<br>The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value:<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| max_connections_per_share | integer | Maximum number of tree connections on share. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each<br>search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a CIFS share

POST `/protocols/cifs/shares`

**Introduced In:** 9.6

Creates a CIFS share.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS share.
- `name` - Name of the CIFS share.
- `path` - Path in the owning SVM namespace that is shared through this share.

#### Recommended optional properties

- `comment` - Optionally choose to add a text comment of up to 256 characters about the CIFS share.
- `acls` - Optionally choose to add share permissions that users and groups have on the CIFS share.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `home_directory` - *false*
- `oplocks` - *true*
- `access_based_enumeration` - *false*
- `change_notify` - *true*
- `encryption` - *false*
- `unix_symlink` - *local*

#### Related ONTAP commands

- `vserver cifs share create`
- `vserver cifs share properties add`
- `vserver cifs share access-control create`

#### Learn more

- DOC /protocols/cifs/shares

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |

| Name | Type | Description |
| --- | --- | --- |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not<br>able to access this share. |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group<br>(also called the "force-group"). The "force-group" must be a predefined group in the UNIX group<br>database. This setting has no effect unless the security style of the volume is UNIX or mixed<br>security style. |

| Name | Type | Description |
| --- | --- | --- |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user. The ONTAP CIFS home directory feature enable us to configure a share that maps to<br>different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| max_connections_per_share | integer | Maximum number of tree connections on share. |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name)<br>variables in any combination with this parameter to generate shares dynamically. |

| Name | Type | Description |
| --- | --- | --- |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share.<br>"Oplocks" allow clients to lock files and cache content locally,<br>which can increase performance for file operations. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern<br>%w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients.<br>The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |

| Name | Type | Description |
| --- | --- | --- |
| vscan_profile | string | Vscan File-Operations Profile<br>The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

### Example request

```
{
  "acls": [
    {
      "permission": "string",
      "type": "string",
      "user_or_group": "ENGDOMAIN\\ad_user"
    }
  ],
  "comment": "HR Department Share",
  "dir_umask": "21",
  "file_umask": "21",
  "force_group_for_create": "string",
  "name": "HR_SHARE",
  "offline_files": "string",
  "path": "/volume_1/eng_vol/",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_symlink": "string",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "vscan_profile": "string"
}
```

### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655628 | CIFS Share Creation with property 'SMB_ENCRYPTION' failed because the CIFS server does not support SMB3.0 |
| 655551 | CIFS Share Creation failed because the specified path does not exist |
| 655577 | The CIFS share name cannot be more than 80 characters long |
| 655399 | Failed to create CIFS share. The CIFS server does not exist for specified SVM |
| 656422 | Failed to create the home directory share because the directory shares must specify a path relative to one or more home directory search paths |
| 656423 | Failed to create CIFS share. The Shares must define an absolute share path |
| 656424 | Failed to create CIFS the administrator share 'c$' because you are not permitted to created any admin shares |
| 656486 | Failed to create CIFS share because the share cannot be made continuously available for FlexCache volumes. |
| 655625 | Failed to create CIFS share. The Shares path is not a valid file-type for CIFS share |
| 656426 | CIFS Share Creation failed because the share name is invalid |
| 655655 | no-strict-security should be set to true only if unix_symlink is configured as "local" or "widelink" |
| 655394 | Failed to create CIFS share because share cannot be made continuously available unless running SMB3 or later. |
| 4849678 | Failed to create CIFS share because the specified UNIX group does not exist |
| 655622 | Invalid value for parameter {max-connections-per-share}. Maximum connections on CIFS share {name} must be between 1 to 4294967295. |
| 655446 | Failed to resolve the security identifier (SID) for the account named {account_name}. Reason: {Reason for failure}. |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |

| Name | Type | Description |
| --- | --- | --- |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |

| Name | Type | Description |
| --- | --- | --- |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group (also called the "force-group"). The "force-group" must be a predefined group in the UNIX group database. This setting has no effect unless the security style of the volume is UNIX or mixed security style. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic. ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user. The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user, a single share with a home directory parameters can be created. In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting %w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| max_connections_per_share | integer | Maximum number of tree connections on share. |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as<br>%w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |

| Name | Type | Description |
| --- | --- | --- |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each<br>search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a CIFS share

DELETE `/protocols/cifs/shares/{svm.uuid}/{name}`

**Introduced In:** 9.6

Deletes a CIFS share.

#### Related ONTAP commands

- `vserver cifs share delete`

#### Learn more

- DOC /protocols/cifs/shares

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Share Name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655393 | Standard admin shares cannot be removed |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a CIFS share

GET `/protocols/cifs/shares/{svm.uuid}/{name}`

**Introduced In:** 9.6

Retrieves a CIFS share.

#### Related ONTAP commands

- `vserver cifs share show`
- `vserver cifs share properties show`

#### Learn more

- DOC /protocols/cifs/shares

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Share Name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| acls | array[acls] | |

| Name | Type | Description |
| --- | --- | --- |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not<br>able to access this share. |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group (also called the "force-group"). The "force-group" must be a predefined group in the UNIX group database. This setting has no effect unless the security style of the volume is UNIX or mixed security style. |
| home_directory | boolean | Specifies whether or not the share is a home directory share, where the share and path names are dynamic.<br>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user. The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user,<br>a single share with a home directory parameters can be created.<br>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting<br>%w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.<br>• Default value: 1<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| max_connections_per_share | integer | Maximum number of tree connections on share. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the CIFS share that you want to create. If this<br>is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name)<br>variables in any combination with this parameter to generate shares dynamically. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally,<br>which can increase performance for file operations. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern<br>%w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients.<br>The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |

| Name | Type | Description |
| --- | --- | --- |
| vscan_profile | string | Vscan File-Operations Profile<br>The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "acls": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "permission": "string",
      "type": "string",
      "user_or_group": "ENGDOMAIN\\ad_user",
      "win_sid_unix_id": "string"
    }
  ],
  "comment": "HR Department Share",
  "dir_umask": "21",
  "file_umask": "21",
  "force_group_for_create": "string",
  "name": "HR_SHARE",
  "offline_files": "string",
  "path": "/volume_1/eng_vol/",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_symlink": "string",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
  },
  "vscan_profile": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

| Name | Type | Description |
| --- | --- | --- |
| win_sid_unix_id | string | Windows SID/UNIX ID depending on access-control type. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a CIFS share

PATCH `/protocols/cifs/shares/{svm.uuid}/{name}`

**Introduced In:** 9.6

Updates a CIFS share.

#### Related ONTAP commands

- `vserver cifs share modify`
- `vserver cifs share properties add`
- `vserver cifs share properties remove`

#### Learn more

- DOC /protocols/cifs/shares

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Share Name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group (also called the "force-group"). The "force-group" must be a predefined group in the UNIX group database. This setting has no effect unless the security style of the volume is UNIX or mixed security style. |
| max_connections_per_share | integer | Maximum number of tree connections on share. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |

| Name | Type | Description |
| --- | --- | --- |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share.<br>"Oplocks" allow clients to lock files and cache content locally,<br>which can increase performance for file operations. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern<br>%w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients.<br>The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |

| Name | Type | Description |
| --- | --- | --- |
| vscan_profile | string | Vscan File-Operations Profile<br>The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

**Example request**

```
{
  "comment": "HR Department Share",
  "dir_umask": "21",
  "file_umask": "21",
  "force_group_for_create": "string",
  "home_directory": true,
  "offline_files": "string",
  "path": "/volume_1/eng_vol/",
  "unix_symlink": "string",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "vscan_profile": "string"
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655628 | 'SMB_ENCRYPTION' property cannot be set on CIFS share because the CIFS server does not support SMB3.0 |
| 655551 | CIFS Share modification failed because the specified path does not exist |
| 655620 | Cannot set symlink properties for admin shares |
| 656420 | Cannot modify the standard share ipc$ |
| 656421 | Cannot modify the standard share admin$ |
| 656422 | Failed to modify the home directory share because the directory shares must specify a path relative to one or more home directory search paths |
| 656423 | Failed to modify CIFS share. The Shares must define an absolute share path |
| 656425 | Failed to modify the CIFS share because the path for an administrative share cannot be modified |
| 656486 | Failed to modify CIFS share because the share cannot be made continuously available for FlexCache volumes. |
| 655395 | Failed to modify the CIFS share because share cannot be made continuously available unless running SMB3 or later. |
| 4849678 | Failed to modify the CIFS share because the specified UNIX group does not exist |
| 655622 | Invalid value for parameter {max-connections-per-share}. Maximum connections on CIFS share {name} must be between 1 to 4294967295. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

acls

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB,
a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.
When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

| Name | Type | Description |
| --- | --- | --- |
| access_based_enumeration | boolean | If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents<br>the display of folders or other shared resources that the user does not have access to. |
| allow_unencrypted_access | boolean | Specifies whether or not the SMB2 clients are allowed to access the encrypted share. |
| attribute_cache | boolean | Specifies that connections through this share cache attributes for a short time to improve performance. |
| browsable | boolean | Specifies whether or not the Windows clients can browse the share. |

| Name | Type | Description |
| --- | --- | --- |
| change_notify | boolean | Specifies whether CIFS clients can request for change notifications for directories on this share. |
| comment | string | Specify the CIFS share descriptions. |
| continuously_available | boolean | Specifies whether or not the clients connecting to this share can open files in a persistent manner.<br>Files opened in this way are protected from disruptive events, such as, failover and giveback.<br>If the Vscan ONTAP feature is used, it is not supported in continuous availability (CA) shares.<br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |
| dir_umask | string | Directory Mode Creation Mask to be viewed as an octal number. |
| encryption | boolean | Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share. |
| file_umask | string | File Mode Creation Mask to be viewed as an octal number. |
| force_group_for_create | string | Specifies that all files that CIFS users create in a specific share belong to the same group (also called the "force-group").<br>The "force-group" must be a predefined group in the UNIX group<br>database. This setting has no effect unless the security style of the volume is UNIX or mixed security style. |

| Name | Type | Description |
| --- | --- | --- |
| max_connections_per_share | integer | Maximum number of tree connections on share. |
| namespace_caching | boolean | Specifies whether or not the SMB clients connecting to this share can cache the directory enumeration<br>results returned by the CIFS servers. |
| no_strict_security | boolean | Specifies whether or not CIFS clients can follow a unix symlinks outside the share boundaries. |
| offline_files | string | Offline Files<br>The supported values are:<br>• none - Clients are not permitted to cache files for offline access.<br>• manual - Clients may cache files that are explicitly selected by the user for offline access.<br>• documents - Clients may automatically cache files that are used by the user for offline access.<br>• programs - Clients may automatically cache files that are used by the user for offline access<br>and may use those files in an offline mode even if the share is available. |
| oplocks | boolean | Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | The fully-qualified pathname in the owning SVM namespace that is shared through this share.<br>If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.<br>ONTAP generates the path dynamically for the connected user and this path is appended to each<br>search path to find the full Home Directory path.<br>• example: /volume_1/eng_vol/<br>• maxLength: 256<br>• minLength: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| show_previous_versions | boolean | Specifies that the previous version can be viewed and restored from the client. |
| show_snapshot | boolean | Specifies whether or not the snapshots can be viewed and traversed by clients. |
| unix_symlink | string | Controls the access of UNIX symbolic links to CIFS clients. The supported values are:<br>• local - Enables only local symbolic links which is within the same CIFS share.<br>• widelink - Enables both local symlinks and widelinks.<br>• disable - Disables local symlinks and widelinks. |
| volume | volume | |

| Name | Type | Description |
| --- | --- | --- |
| vscan_profile | string | Vscan File-Operations Profile The supported values are:<br>• no_scan - Virus scans are never triggered for accesses to this share.<br>• standard - Virus scans can be triggered by open, close, and rename operations.<br>• strict - Virus scans can be triggered by open, read, close, and rename operations.<br>• writes_only - Virus scans can be triggered only when a file that has been modified is closed. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage share-level ACL

### Manage share-level ACL

#### Overview

Access to files and folders can be secured over a network by configuring share access control lists (ACLs) on CIFS shares. Share-level ACLs can be configured by using either Windows users and groups or UNIX users and groups. A share-level ACL consists of a list of access control entries (ACEs). Each ACE contains a user or group name and a set of permissions that determines user or group access to the share, regardless of the security style of the volume or qtree containing the share.

When an SMB user tries to access a share, ONTAP checks the share-level ACL to determine whether access should be granted. A share-level ACL only restricts access to files in the share; it never grants more access than the file level ACLs.

#### Examples

### Creating a CIFS share ACL

To create a share ACL for a CIFS share, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/shares{svm.uuid}/{share}/acls

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-
11e8-a96e-0050568ea3cb/sh1/acls?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"permission\": \"no_access\", \"type\": \"windows\", \"user_or_group\":
\"root\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "user_or_group": "root",
    "type": "windows",
    "permission": "no_access",
    "sid": "S-1-1-0"
  }
]
}
```

### Retrieving all CIFS shares ACLs for a specific CIFS share for a specific SVM in the cluster

```
# The API:
GET /api/protocols/cifs/shares/{svm.uuid}/{share}/acls

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-
11e8-a96e-
0050568ea3cb/sh1/acls?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs1"
    },
    "share": "sh1",
    "user_or_group": "Everyone",
    "type": "windows",
    "permission": "full_control",
    "sid": "S-1-1-0"
  },
  {
    "svm": {
      "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs1"
    },
    "share": "sh1",
    "user_or_group": "root",
    "type": "windows",
    "permission": "no_access",
    "sid": "S-1-1-0"
  }
],
"num_records": 2
}
```

#### Retrieving a CIFS share ACLs for a user or a group of type Windows or type UNIX on a CIFS share for a specific SVM

```
# The API:
GET
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-
11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
  "name": "vs1"
},
"share": "sh1",
"user_or_group": "everyone",
"type": "windows",
"permission": "full_control",
"sid": "S-1-1-0"
}
```

#### Updating a CIFS share ACLs of a user or group on a CIFS share for a specific SVM

The CIFS share ACL being modified is identified by the UUID of its SVM, the CIFS share name, user or group name and the type of the user or group.

```
# The API:
PATCH
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-
11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"permission\": \"no_access\"}"
```

#### Removing a CIFS share ACLs of a user or group on a CIFS Share for a specific SVM

The CIFS share ACL being removed is identified by the UUID of its SVM, the CIFS share name, user or group name and the type of the user or group.

```
# The API:
DELETE
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-
11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept:
application/json"
```

---

### Retrieve a share-level ACL on a CIFS share

GET `/protocols/cifs/shares/{svm.uuid}/{share}/acls`

**Introduced In:** 9.6

Retrieves the share-level ACL on a CIFS share.

#### Related ONTAP commands

- `vserver cifs share access-control show`

#### Learn more

- DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share | string | path | True | CIFS Share Name |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.9 |
| sid | string | query | False | Filter by sid<br>• Introduced in: 9.13 |
| type | string | query | False | Filter by type |
| permission | string | query | False | Filter by permission |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_id | integer | query | False | Filter by unix_id<br>• Introduced in: 9.16 |
| user_or_group | string | query | False | Filter by user_or_group |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_share_acl] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "permission": "string",
      "share": "string",
      "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "unix_id": 100,
      "user_or_group": "ENGDOMAIN\\ad_user"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_share_acl

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |
| sid | string | Specifies the user or group secure identifier (SID). |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a share-level ACL on a CIFS share

POST `/protocols/cifs/shares/{svm.uuid}/{share}/acls`

**Introduced In:** 9.6

Creates a share-level ACL on a CIFS share.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the share acl.
- `share` - Existing CIFS share in which to create the share acl.
- `user_or_group` - Existing user or group name for which the acl is added on the CIFS share.
- `permission` - Access rights that a user or group has on the defined CIFS share.

#### Default property values

- `type` - *windows*

#### Related ONTAP commands

- `vserver cifs share access-control create`

#### Learn more

- DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share | string | path | True | CIFS Share Name |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |
| sid | string | Specifies the user or group secure identifier (SID). |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

### Example request

```
{
  "permission": "string",
  "share": "string",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "unix_id": 100,
  "user_or_group": "ENGDOMAIN\\ad_user"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655446 | Failed to resolve the security identifier (SID) for the account named {user_or_group}. Reason: {Reason}. |
| 4849678 | Failed to resolve {user_or_group} name to a UNIX ID. Reason: {Reason}. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_share_acl

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |
| sid | string | Specifies the user or group secure identifier (SID). |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the type of the user or group to add to the access control<br>list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a share-level ACL on a CIFS share

DELETE `/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}`

**Introduced In:** 9.6

Deletes a share-level ACL on a CIFS share.

#### Related ONTAP commands

- `vserver cifs share access-control delete`

#### Learn more

- DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share | string | path | True | Share name |
| user_or_group | string | path | True | User or group name |
| type | string | path | True | User or group type |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a share-level ACL on a CIFS share for a user or group

GET `/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}`

**Introduced In:** 9.6

Retrieves the share-level ACL on CIFS share for a specified user or group.

#### Related ONTAP commands

- `vserver cifs share access-control show`

#### Learn more

- DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share | string | path | True | Share name |
| user_or_group | string | path | True | User or group name |
| type | string | path | True | User or group type |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |
| sid | string | Specifies the user or group secure identifier (SID). |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed:<br>• windows - Windows user or group<br>• unix_user - UNIX user<br>• unix_group - UNIX group |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |
| user_or_group | string | Specifies the user or group name to add to the access control list of a CIFS share. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "permission": "string",
  "share": "string",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "unix_id": 100,
  "user_or_group": "ENGDOMAIN\\ad_user"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a share-level ACL on a CIFS share

PATCH `/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}`

**Introduced In:** 9.6

Updates a share-level ACL on a CIFS share.

#### Related ONTAP commands

- `vserver cifs share access-control modify`

#### Learn more

- DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| share | string | path | True | Share name |
| user_or_group | string | path | True | User or group name |
| type | string | path | True | User or group type |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |

| Name | Type | Description |
| --- | --- | --- |
| sid | string | Specifies the user or group secure identifier (SID). |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |

**Example request**

```
{
  "permission": "string",
  "share": "string",
  "sid": "S-1-5-21-256008430-3394229847-3930036330-1001",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_id": 100
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 655516 | The share ACL does not exist for given user and share |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_share_acl

The permissions that users and groups have on a CIFS share.

| Name | Type | Description |
|---|---|---|
| permission | string | Specifies the access rights that a user or group has on the defined CIFS Share.<br>The following values are allowed:<br>• no_access - User does not have CIFS share access<br>• read - User has only read access<br>• change - User has change access<br>• full_control - User has full_control access |
| share | string | CIFS share name |
| sid | string | Specifies the user or group secure identifier (SID). |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_id | integer | Specifies the UNIX user or group identifier (UID/GID). |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage CIFS UNIX symlink mapping

### Manage CIFS UNIX symlink mapping

#### Overview

ONTAP allows both CIFS and NFS to access the same datastore. This datastore can contain symbolic links which are files, created by UNIX clients. It contains a reference to another file or directory. If an SMB client accesses a symbolic link, it is redirected to the target file or directory that the symbolic link refers to. The symbolic links can point to files within the volume that contain the share, or to files that are contained in other volumes on the Storage Virtual Machine (SVM), or even to volumes contained on other SVMs.

There are two types of symbolic links:

---

**Relative** A relative symbolic link contains a reference to the file or directory relative to its parent directory. Therefore, the path of the file it is referring to should not begin with a backslash (/). If you enable symbolic links on a share, relative symbolic links work without UNIX symlink mapping.

**Absolute** An absolute symbolic link contains a reference to a file or directory in the form of an absolute path. Therefore, the path of the file it is referring to should begin with a backslash (/). An absolute symbolic link can refer to a file or directory within or outside of the file system of the symbolic link. If the target is not in the same local file system, the symbolic link is called a "widelink". If the symbolic link is enabled on a share and absolute symbolic links do not work right away, the mapping between the UNIX path of the symbolic link to the destination CIFS path must be created. When creating absolute symbolic link mappings, locality could be either "local" or "widelink" and it must be specified. If UNIX symlink mapping is created for a file or directory which is outside of the local share but the locality is set to "local", ONTAP does not allow access to the target.

A UNIX symbolic link support could be added to SMB shares by specifying the *unix_symlink* property during the creation of SMB shares or at any time by modifying the existing SMB *unix_symlink* property. UNIX symbolic link support is enabled by default.

#### Examples

### Creating a UNIX symlink mapping for CIFS shares

To create UNIX symlink mappings for SMB shares, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/unix-symlink-mapping

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping?return_records=true" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"svm\": { \"name\": \"vs1\", \"uuid\":
\"000c5cd2-ebdf-11e8-a96e-0050568ea3cb\" }, \"target\": {
\"home_directory\": false, \"locality\": \"local\", \"path\":
\"/dir1/dir2/\", \"server\": \"cifs123\", \"share\": \"sh1\" },
\"unix_path\": \"/mnt/eng_volume/\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs1"
    },
    "unix_path": "/mnt/eng_volume/",
    "target": {
      "share": "sh1",
      "path": "/dir1/dir2/",
      "server": "cifs123",
      "locality": "local",
      "home_directory": false
    }
  }
]
}
```

#### Retrieving UNIX symlink mappings for all SVMs in the cluster

```
# The API:
GET /api/protocols/cifs/unix-symlink-mapping

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping?fields=*&return_records=true&return_timeout=15" -H "accept:
application/hal+json"
```

```
# The response:
{
"records": [
  {
    "svm": {
      "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
        }
      }
    },
    "unix_path": "/mnt/eng_volume/",
    "target": {
      "share": "sh1",
      "path": "/dir1/dir2/",
      "server": "CIFS123",
      "locality": "local",
      "home_directory": false
    },
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/unix-symlink-mapping/000c5cd2-ebdf-
11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F"
      }
    }
  },
  {
    "svm": {
      "uuid": "1d30d1b1-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/1d30d1b1-ebdf-11e8-a96e-0050568ea3cb"
        }
      }
    },
    "unix_path": "/mnt/eng_volume/",
    "target": {
      "share": "ENG_SHARE",
      "path": "/dir1/dir2/",
      "server": "ENGCIFS",
      "locality": "widelink",
      "home_directory": false
    },
```

```
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/unix-symlink-mapping/1d30d1b1-ebdf-
11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/cifs/unix-symlink-
mapping?fields=*&return_records=true&return_timeout=15"
  }
}
}
```


#### Retrieving a specific UNIX symlink mapping for an SVM

The mapping being returned is identified by the UUID of its SVM and the unix-path.


```
# The API:
GET /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H
"accept: application/json"

# The response:
{
"svm": {
  "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
 "name": "vs1"
},
"unix_path": "/mnt/eng_volume/",
"target": {
  "share": "sh1",
  "path": "/dir1/dir2/",
  "server": "CIFS123",
  "locality": "local",
  "home_directory": false
}
}
```

#### Updating a specific UNIX symlink mapping for an SVM

The mapping being modified is identified by the UUID of its SVM and the unix-path.

```
# The API:
PATCH /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"target\": { \"home_directory\": true, \"locality\": \"widelink\",
\"path\": \"/new_path/\", \"server\": \"HR_SERVER\", \"share\": \"sh2\"
}}"
```

#### Removing a specific UNIX symlink mapping for an SVM

The mapping being removed is identified by the UUID of its SVM and the unix-path.

```
# The API:
DELETE /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H
"accept: application/json"
```

### Retrieve UNIX symbolic link mappings for CIFS clients

GET `/protocols/cifs/unix-symlink-mapping`

**Introduced In:** 9.6

Retrieves UNIX symbolic link mappings for CIFS clients.

#### Related ONTAP commands

- `vserver cifs symlink show`

#### Learn more

- DOC /protocols/cifs/unix-symlink-mapping

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| target.home_directory | boolean | query | False | Filter by target.home_directory |
| target.locality | string | query | False | Filter by target.locality |
| target.path | string | query | False | Filter by target.path<br>• maxLength: 256 |
| target.server | string | query | False | Filter by target.server<br>• maxLength: 45 |
| target.share | string | query | False | Filter by target.share<br>• maxLength: 80 |
| unix_path | string | query | False | Filter by unix_path<br>• maxLength: 256 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[cifs_symlink_mapping] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "locality": "string",
        "path": "/dir1/dir2/",
        "server": "ENGCIFS",
        "share": "ENG_SHARE"
      },
      "unix_path": "/mnt/eng_volume/"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_target

| Name | Type | Description |
|---|---|---|
| home_directory | boolean | Specify if the destination share is a home directory. |

| Name | Type | Description |
| --- | --- | --- |
| locality | string | Specifies whether the CIFS symbolic link is a local link or wide link.<br>The following values are supported:<br>• local - Local symbolic link maps only to the same CIFS share.<br>• widelink - Wide symbolic link maps to any CIFS share on the network. |
| path | string | Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted. |
| server | string | Specifies the destination CIFS server where the<br>UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'.<br>You can specify the value in any of the following formats:<br>• DNS name of the CIFS server.<br>• IP address of the CIFS server.<br>• NetBIOS name of the CIFS server. |
| share | string | Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing. |

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |
| unix_path | string | Specifies the UNIX path prefix to be matched for the mapping. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a UNIX symbolic link mapping for a CIFS client**

POST `/protocols/cifs/unix-symlink-mapping`

**Introduced In:** 9.6

Creates a UNIX symbolic link mapping for a CIFS client.

**Required properties**

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS unix-symlink-mapping.

- `unix_path` - UNIX path to which the CIFS symlink mapping to be created.
- `target.share` - CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.
- `target.path` - CIFS path on the destination to which the symbolic link maps.

#### Default property values

- `target.server` - *Local_NetBIOS_Server_Name*
- `locality` - *local*
- `home_directory` - *false*

#### Related ONTAP commands

- `vserver cifs symlink create`

#### Learn more

- DOC /protocols/cifs/unix-symlink-mapping

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |
| unix_path | string | Specifies the UNIX path prefix to be matched for the mapping. |

### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[cifs_symlink_mapping] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "locality": "string",
        "path": "/dir1/dir2/",
        "server": "ENGCIFS",
        "share": "ENG_SHARE"
      },
      "unix_path": "/mnt/eng_volume/"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655654 | Must specify the target CIFS share while creating path mapping entries with localities "local" or "widelink" |
| 655572 | The target path contains illegal characters or is too long |
| 655574 | The target server contains illegal characters or is too long |
| 655436 | If the locality is "local", the target server must be blank or must match the CIFS NetBIOS name for given SVM |

| Error Code | Description |
| --- | --- |
| 655439 | The Specified target server is local CIFS server for given SVM but the locality is specified as "widelink" |
| 655546 | Failed to create symlink mapping because administrative share cannot be used as target share |
| 655437 | Failed to create the symlink mapping with locality "local" because the target share does not exist for the specified SVM |
| 655429 | UNIX path must begin and end with a "/" |
| 655430 | Target path must begin and end with a "/" |
| 655399 | Failed to get the CIFS server for specified SVM |
| 656481 | Failed to create the widelink because the target share does not exist for the specified SVM |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_target

| Name | Type | Description |
|---|---|---|
| home_directory | boolean | Specify if the destination share is a home directory. |
| locality | string | Specifies whether the CIFS symbolic link is a local link or wide link.<br>The following values are supported:<br>• local - Local symbolic link maps only to the same CIFS share.<br>• widelink - Wide symbolic link maps to any CIFS share on the network. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted. |
| server | string | Specifies the destination CIFS server where the<br>UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'.<br>You can specify the value in any of the following formats:<br>• DNS name of the CIFS server.<br>• IP address of the CIFS server.<br>• NetBIOS name of the CIFS server. |
| share | string | Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing. |

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |

| Name | Type | Description |
|---|---|---|
| unix_path | string | Specifies the UNIX path prefix to be matched for the mapping. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a UNIX symbolic link mapping for CIFS clients

DELETE `/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}`

**Introduced In:** 9.6

Deletes the UNIX symbolic link mapping for CIFS clients.

#### Related ONTAP commands

- `vserver cifs symlink delete`

#### Learn more

- DOC /protocols/cifs/unix-symlink-mapping

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_path | string | path | True | UNIX symbolic link path |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a UNIX symbolic link mapping for CIFS clients

GET `/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}`

**Introduced In:** 9.6

Retrieves a UNIX symbolic link mapping for CIFS clients.

#### Related ONTAP commands

- `vserver cifs symlink show`

#### Learn more

- DOC /protocols/cifs/unix-symlink-mapping

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_path | string | path | True | UNIX symbolic link path |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |
| unix_path | string | Specifies the UNIX path prefix to be matched for the mapping. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_target

| Name | Type | Description |
|---|---|---|
| home_directory | boolean | Specify if the destination share is a home directory. |
| locality | string | Specifies whether the CIFS symbolic link is a local link or wide link.<br>The following values are supported:<br>• local - Local symbolic link maps only to the same CIFS share.<br>• widelink - Wide symbolic link maps to any CIFS share on the network. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted. |
| server | string | Specifies the destination CIFS server where the<br>UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'.<br>You can specify the value in any of the following formats:<br>• DNS name of the CIFS server.<br>• IP address of the CIFS server.<br>• NetBIOS name of the CIFS server. |
| share | string | Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a UNIX symbolic link mapping for CIFS clients

PATCH `/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}`

**Introduced In:** 9.6

Updates the UNIX symbolic link mapping for CIFS clients.

#### Related ONTAP commands

- `vserver cifs symlink modify`

#### Learn more

- DOC /protocols/cifs/unix-symlink-mapping

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_path | string | path | True | UNIX symbolic link path |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |

#### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655437 | Failed to modify the symlink mapping with locality "local" because the target share does not exist for the specified SVM |
| 655573 | Failed to modify the symlink mapping to target path because it contains illegal characters or is too long |
| 655575 | Failed to modify the symlink mapping to target server because it contains illegal characters or is too long |
| 655547 | Failed to modify symlink mapping because administrative share cannot be used as target share |
| 656481 | Failed to modify the widelink because the target share does not exist for the specified SVM |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cifs_target

| Name | Type | Description |
|---|---|---|
| home_directory | boolean | Specify if the destination share is a home directory. |
| locality | string | Specifies whether the CIFS symbolic link is a local link or wide link.<br>The following values are supported:<br>• local - Local symbolic link maps only to the same CIFS share.<br>• widelink - Wide symbolic link maps to any CIFS share on the network. |

| Name | Type | Description |
| --- | --- | --- |
| path | string | Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted. |
| server | string | Specifies the destination CIFS server where the<br>UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'.<br>You can specify the value in any of the following formats:<br>• DNS name of the CIFS server.<br>• IP address of the CIFS server.<br>• NetBIOS name of the CIFS server. |
| share | string | Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing. |

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | cifs_target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Perform bulk import for CIFS user groups

### Perform bulk import for CIFS user groups

#### Overview

Bulk import of the CIFS local users, groups and group membership information can be done from the specified Uniform Resource Identifier (URI). This replaces the existing contents of the CIFS local users, groups and group memberships. This API is used to bulk import from the specified URI, get the status of the last import and to upload the import status to the specified URI.

#### Retrieving import status of the last bulk import

The bulk-import GET endpoint retrieves the status of the last bulk-import operation of the specified SVM.

#### Examples

### Retrieving the status of a successful bulk import

```
# The API:
/api/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/users-and-groups/bulk-
import/6de1d39d-1473-11ec-b0cf-0050568e4acc" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "6de1d39d-1473-11ec-b0cf-0050568e4acc",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/6de1d39d-1473-11ec-b0cf-0050568e4acc"
    }
  }
},
"import_uri": {
  "path": "http://<import_uri>/file.7z"
},
"state": "success",
"detailed_status": {
  "message": "Operation completed successfully.",
  "code": "0"
},
"elements_ignored": 0,
"elements_imported": 2,
"_links": {
  "self": {
    "href": "/api/protocols/cifs/users-and-groups/import/6de1d39d-1473-
11ec-b0cf-0050568e4acc"
  }
}
}
```

**Retrieving the status of a bulk import that failed**

```
# The API:
/api/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/users-and-groups/bulk-
import/6de1d39d-1473-11ec-b0cf-0050568e4acc" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "6de1d39d-1473-11ec-b0cf-0050568e4acc",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/6de1d39d-1473-11ec-b0cf-0050568e4acc"
    }
  }
},
"import_uri": {
  "path": "http://<import_uri>/file.7z"
},
"state": "success",
"detailed_status": {
  "message": "Failed parsing line 1 of the input file. Check syntax and
contents.",
  "code": "655698"
},
"elements_ignored": 0,
"elements_imported": 0,
"_links": {
  "self": {
    "href": "/api/protocols/cifs/users-and-groups/import/6de1d39d-1473-
11ec-b0cf-0050568e4acc"
  }
}
}
```

**Retrieving bulk import information for CIFS local users, groups, and group memberships**

```
# The API:
/api/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/users-and-groups/bulk-
import/6de1d39d-1473-11ec-b0cf-0050568e4acc" -H "accept: application/json"
-d
'{"import_uri.username":"user1","import_uri.password":"aaaa","decryption_p
assword":"cccc","import_uri.path":"http://example.com/file1.7z"}'

# The response:
{
"job": {
  "uuid": "e8293262-1af0-11ec-8243-0050568e4acc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e8293262-1af0-11ec-8243-0050568e4acc"
    }
  }
}
}
```

**Retrieving status upload information of the last import operation for the specified URI**

```
# The API:
/api/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/users-and-groups/bulk-
import/6de1d39d-1473-11ec-b0cf-0050568e4acc" -H "accept: application/json"
-d
'{"status_uri.username":"user1","status_uri.password":"aaaa","status_uri.p
ath":"http://example.com/fileupload.7z"}'

# The response:
{
"job": {
  "uuid": "e8293262-1af0-11ec-8243-0050568e4acd",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e8293262-1af0-11ec-8243-0050568e4acd"
    }
  }
}
}
```

---

### Retrieve import operation status for CIFS local users, groups, and group memberships

GET `/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}`

**Introduced In:** 9.11

Retrieves information about the import operation status of the CIFS local users, groups, and group memberships.

#### Related ONTAP commands

- `vserver cifs users-and-groups import get-status`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| detailed_status | detailed_status | |
| elements_ignored | integer | Number of elements ignored. |
| elements_imported | integer | Number of elements successfully imported. |
| import_uri | import_uri | |
| state | string | Operation status. |
| status_uri | status_uri | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "detailed_status": {
    "code": "6684732",
    "message": "string"
  },
  "elements_ignored": 0,
  "elements_imported": 0,
  "import_uri": {
    "path": "http://web.sample.com/web1/file1.7z"
  },
  "state": "string",
  "status_uri": {
    "path": "http://web.sample.com/web1/file1.7z"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

detailed_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the import status failure. |
| message | string | Detailed description of the import status. |

import_uri

| Name | Type | Description |
|---|---|---|
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |

status_uri

| Name | Type | Description |
|---|---|---|
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Upload the bulk-import status of an SVM to a URI

PATCH `/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}`

**Introduced In:** 9.11

Upload the status of the bulk-import of the specified SVM to the specified URI.

#### Important notes

- Only the status of the last bulk-import will be uploaded and not the status of the previous bulk-imports.

#### Required properties

- status_uri.path - URI to which the status needs to be uploaded.

#### Optional properties

- status_uri.username - Username of the specified URI.
- status_uri.password - Password of the specified URI.

#### Related ONTAP commands

- `vserver cifs users-and-groups import get-status`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| elements_ignored | integer | Number of elements ignored. |
| elements_imported | integer | Number of elements successfully imported. |
| state | string | Operation status. |
| status_uri | status_uri | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "elements_ignored": 0,
  "elements_imported": 0,
  "state": "string",
  "status_uri": {
    "password": "string",
    "path": "http://web.sample.com/web1/file1.7z",
    "username": "user1"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default, Error
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

detailed_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the import status failure. |
| message | string | Detailed description of the import status. |

import_uri

| Name | Type | Description |
|---|---|---|
| password | string | Password of the specified URI. |
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |
| username | string | Username of the specified URI. |

status_uri

| Name | Type | Description |
|---|---|---|
| password | string | Password of the specified URI. |
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |
| username | string | Username of the specified URI. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_users_and_groups_import

| Name | Type | Description |
|---|---|---|
| elements_ignored | integer | Number of elements ignored. |
| elements_imported | integer | Number of elements successfully imported. |
| state | string | Operation status. |
| status_uri | status_uri | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Load CIFS local users, groups and group membership files from a URL

POST `/protocols/cifs/users-and-groups/bulk-import/{svm.uuid}`

**Introduced In:** 9.11

Loads CIFS local users,groups and group memberships file from the specified URL.

#### Important notes

Existing CIFS local users, groups, and group memberships will be replaced with the contents of the file.

#### Required properties

- import_uri.path
- decryption_password

#### Optional properties

- import_uri.username
- import_uri.password

#### Related ONTAP commands

- `vserver cifs users-and-groups import load-from-uri`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| decryption_password | string | Password to decrypt the compressed file. |
| elements_ignored | integer | Number of elements ignored. |
| elements_imported | integer | Number of elements successfully imported. |

| Name | Type | Description |
| --- | --- | --- |
| import_uri | import_uri | |
| state | string | Operation status. |
| status_uri | status_uri | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "decryption_password": "string",
  "elements_ignored": 0,
  "elements_imported": 0,
  "import_uri": {
    "password": "string",
    "path": "http://web.sample.com/web1/file1.7z",
    "username": "user1"
  },
  "state": "string",
  "status_uri": {
    "password": "string",
    "path": "http://web.sample.com/web1/file1.7z",
    "username": "user1"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655688 | Invalid file decryption password. Valid password characters are: a-z, A-Z, and 0-9. |
| 655660 | This operation is only permitted on data SVMs. |
| 655661 | URI username, URI password and the file decryption password cannot exceed 128 characters. |
| 655689 | File decryption password cannot be empty. |
| 655399 | CIFS server does not exist. |
| 655683 | Domain Security Identifier is already present in Active Directory. |
| 655685 | Cluster is busy processing a local users and groups transition task. |
| 655687 | Unknown file schema version. |
| 655690 | Processing the import file failed. Check the file decryption password, compression format and content. |

| Error Code | Description |
| --- | --- |
| 655691 | Error occurred before processing the file. Verify the input parameters. |
| 655692 | Domain name that qualifies the local users and groups must match the CIFS server name. |
| 655693 | All local users and local groups must be prefixed with the NetBIOS domain name. |
| 655693 | Well-known user of group will not be present in this SVM. |
| 655695 | Group membership of the specified user will not be created. |
| 655697 | Failed processing a line of the bulk import file. |
| 655698 | Failed to parse a line of the bulk import file. Check the syntax and contents. |
| 655699 | Invalid number of elements in a line of the bulk import file. |
| 655700 | Failed processing a field index of a line in the input bulk import file. |
| 655701 | Unknown element type. |
| 655702 | Invalid URI. |
| 655703 | Failed to download the import file. Network error occurred. |
| 655709 | Failed to queue CIFS local users and groups import job since a similar job already exists. Retry the operation after some time. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

detailed_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the import status failure. |
| message | string | Detailed description of the import status. |

import_uri

| Name | Type | Description |
|---|---|---|
| password | string | Password of the specified URI. |
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |
| username | string | Username of the specified URI. |

status_uri

| Name | Type | Description |
|---|---|---|
| password | string | Password of the specified URI. |
| path | string | URI from which to load the input file containing the CIFS local users and groups. The file must be encrypted using the 7zip utility. URI can be FTP or HTTP. |
| username | string | Username of the specified URI. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_cifs_users_and_groups_import

| Name | Type | Description |
|---|---|---|
| decryption_password | string | Password to decrypt the compressed file. |
| elements_ignored | integer | Number of elements ignored. |
| elements_imported | integer | Number of elements successfully imported. |
| import_uri | import_uri | |
| state | string | Operation status. |
| status_uri | status_uri | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage CIFS user group privileges

### Manage CIFS user group privileges

#### Overview

Privileges associated with local or Active Directory users or groups defines the permissions for the specified user or group. You can use this API to display and/or control privileges of local or Active Directory users or groups.

#### Retrieving the privileges of a specific local or Active Directory user or group and an SVM

The users and groups privileges GET endpoint retrieves privileges of the specified local or Active Directory user or group and the SVM.

#### Examples

**Retrieving the privileges of all of the users or groups of data SVMs.**

---

```
# The API:
/api/protocols/cifs/users-and-groups/privileges

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/users-and-
groups/privileges?fields=*" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "name": "VS1.CIFS\\user1",
```

```
    "privileges": [
      "sechangenotifyprivilege",
      "setakeownershipprivilege"
    ]
  },
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "name": "ACTIVE_DIRECTORY\\user",
    "privileges": [
      "sebackupprivilege",
      "setakeownershipprivilege"
    ]
  },
  {
    "svm": {
      "uuid": "0ac79c37-3867-11eb-bece-0050568ed0a2",
      "name": "vs2"
    },
    "name": "VS2.CIFS\\group1",
    "privileges": [
      "sesecurityprivilege",
      "sebackupprivilege",
      "serestoreprivilege"
    ]
  }
]
}
```


**Retrieving the privileges of the specific SVM and user or group**

```
# The API:
/api/protocols/cifs/users-and-groups/{svm.uuid}/{name}/privileges

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/users-and-
groups/privileges/25b363a6-2971-11eb-88e1-0050568eefd4/user1" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
  "name": "vs1"
},
"name": "VS1.CIFS\\user1",
"privileges": [
  "sechangenotifyprivilege",
  "setakeownershipprivilege"
]
}
```

#### Adding privileges to the local or Active Directory user or group

The users and groups privileges POST endpoint adds privileges to the specified local or Active Directory user or group and the SVM.

### Adding the privileges to the local user 'user1'

```
# The API:
/api/protocols/cifs/users-and-groups/privileges/{svm.uuid}/{name}

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/users-and-
groups/privileges/179d3c85-7053-11e8-b9b8-005056b41bd1/user1" -H "accept:
application/json" -H "Content-Type: application/json" -d '{ "privileges":
[ "SeSecurityPrivilege", "SeBackupPrivilege", "SeRestorePrivilege"] }'
```

#### Updating the privileges of the local or Active Directory user or group of a specific SVM

#### Example

### Updating the privileges of local user 'user1' in SVM 'vs1' to 'SeRestorePrivilege' and 'SeSecurityPrivilege'

```
# The API:
/api/protocols/cifs/users-and-groups/privileges/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/users-and-
groups/privileges/179d3c85-7053-11e8-b9b8-005056b41bd1/user1" -H "accept:
application/json" -d '{ "privileges": [ "SeRestorePrivilege",
"SeSecurityPrivilege"] }'
```

**Reset all the privileges associated with the local user 'user1' in SVM 'vs1'**

```
# The API:
/api/protocols/cifs/users-and-groups/privileges/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/users-and-
groups/privileges/179d3c85-7053-11e8-b9b8-005056b41bd1/user1" -H "accept:
application/json" -d '{ "privileges": [ ] }'
```

**Retrieve privileges for all local, or Active Directory users or groups and SVMs**

GET `/protocols/cifs/users-and-groups/privileges`

**Introduced In:** 9.10

Retrieves privileges of the specified local or Active Directory user or group and SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups privilege show`

#### Learn more

- DOC /protocols/cifs/users-and-groups/privileges

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name |
| privileges | string | query | False | Filter by privileges |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of local or Active Directory user or group records. |

| Name | Type | Description |
| --- | --- | --- |
| records | array[user_group_privileges] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "user1",
      "privileges": [
        "string"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

user_group_privileges

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Local or Active Directory user or group name. |

| Name | Type | Description |
| --- | --- | --- |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add privileges to a local, or Active Directory user or group and SVM

POST `/protocols/cifs/users-and-groups/privileges`

**Introduced In:** 9.10

Adds privileges to the specified local or Active Directory user or group and SVM.

#### Important note

- Specified privileges are appended to the existing list of privileges.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM for which privileges are added to user or group.
- `name` - Existing local or Active Directory user or group for which privileges are to be added.
- `privileges` - List of privileges to be added to a user or group.

#### Related ONTAP commands

- `vserver cifs users-and-groups privilege add-privilege`

#### Learn more

- DOC /protocols/cifs/users-and-groups/privileges

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| name | string | Local or Active Directory user or group name. |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "name": "user1",
  "privileges": [
    "string"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655660 | The operation is allowed only on data SVMs. |
| 655673 | Failed to resolve the user or group. |
| 655730 | The specified local user to which privileges are to be associated to does not exist. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

user_group_privileges

| Name | Type | Description |
|---|---|---|
| name | string | Local or Active Directory user or group name. |

| Name | Type | Description |
| --- | --- | --- |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve privileges for a local, or Active Directory user or group and SVM

GET `/protocols/cifs/users-and-groups/privileges/{svm.uuid}/{name}`

**Introduced In:** 9.10

Retrieves privileges of the specified local or Active Directory user or group and SVM.

#### Related ONTAP commands

- `vserver cifs users-and-groups privilege show`

#### Learn more

- DOC /protocols/cifs/users-and-groups/privileges

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Local or Active Directory user or group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Local or Active Directory user or group name. |

| Name | Type | Description |
| --- | --- | --- |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "user1",
  "privileges": [
    "string"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update privileges for a local, or Active Directory user or group and SVM

PATCH `/protocols/cifs/users-and-groups/privileges/{svm.uuid}/{name}`

**Introduced In:** 9.10

Updates privileges of the specified local or Active Directory user or group and SVM.

#### Important note

- Specified privileges will replace all the existing privileges associated with the user or group.
- To reset privileges associated with the user or group, specify the privileges list as empty.

#### Related ONTAP commands

- `vserver cifs users-and-groups privilege reset-privilege`

#### Learn more

- DOC /protocols/cifs/users-and-groups/privileges

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Local or Active Directory user or group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |

**Example request**

```
{
  "privileges": [
    "string"
  ]
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 262196 | Field 'svm.name' is not supported in the body of PATCH request. |
| 262203 | Field 'svm.uuid' is not supported in the body of PATCH request. |
| 655673 | Failed to resolve the user or group. |
| 655730 | The specified local user to which privileges are to be associated to does not exist. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

user_group_privileges

| Name | Type | Description |
| --- | --- | --- |
| privileges | array[string] | An array of privileges associated with the local or Active Directory user or group.<br>The available values are:<br>• SeTcbPrivilege - Allows user to act as part of the operating system<br>• SeBackupPrivilege - Allows user to back up files and directories, overriding any ACLs<br>• SeRestorePrivilege - Allows user to restore files and directories, overriding any ACLs<br>• SeTakeOwnershipPrivilege - Allows user to take ownership of files or other objects<br>• SeSecurityPrivilege - Allows user to manage auditing and viewing/dumping/clearing the security log<br>• SeChangeNotifyPrivilege - Allows user to bypass traverse checking |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## View file security permissions

### View file security permissions

#### Overview

This API displays the effective permission granted to a Windows or UNIX user on the specified file or folder path. A path within the FlexCache volume is not supported.

#### Examples

**Retrieving the effective permission for the specified Windows user on the specified path of an SVM.**

```
# The API:
curl -X GET "https://10.63.26.252/api/protocols/file-security/effective-
permissions/cf5f271a-1beb-11ea-8fad-
005056bb645e/administrator/windows/%2F?share.name=sh1&return_records=true"
-H "accept: application/json" -H "authorization: Basic <CREDENTIALS>"

# The response:
{
"svm": {
  "uuid": "cf5f271a-1beb-11ea-8fad-005056bb645e",
  "name": "vs1"
},
"user": "administrator",
"type": "windows",
"path": "/",
"share": {
  "path": "/"
},
"file_permission": [
  "read",
  "write",
  "append",
  "read_ea",
  "write_ea",
  "execute",
  "delete_child",
  "read_attributes",
  "write_attributes",
  "delete",
  "read_control",
  "write_dac",
  "write_owner",
  "synchronize",
  "system_security"
],
"share_permission": [
  "read",
  "read_ea",
  "execute",
  "read_attributes",
  "read_control",
  "synchronize"
]
}
```

### Retrieve effective security permissions on a file

GET `/protocols/file-security/effective-permissions/{svm.uuid}/{path}`

**Introduced In:** 9.8

Retrieves effective security permissions on a file.

#### Related ONTAP commands

- `vserver security file-directory show-effective-permissions`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | File Path |
| user | string | query | True | User_Name |
| share.name | string | query | False | Share Name |
| type | string | query | False | User Type |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| file_permissions | array[string] | Specifies the effective permission granted to a user on the specified file or folder path. |
| path | string | Specifies the path of the file or the folder for which you want to display effective permissions.<br>The path is relative to the SVM root volume. If "-share-name" is specified then path will be relative to the share path. |

| Name | Type | Description |
| --- | --- | --- |
| share | share | |
| share_permissions | array[string] | Specifies the effective permission granted to a user on the specified file or folder path. |
| svm | svm_reference | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the user type. The following values are allowed:<br>• windows - Windows user<br>• unix - UNIX user |
| user | string | Specifies the user for which effective permission needs to be displayed for the specified path. |

### Example response


```
{
  "file_permissions": [
    "string"
  ],
  "path": "/dir1/dir2",
  "share": {
    "name": "string",
    "path": "string"
  },
  "share_permissions": [
    "string"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "user": "cifs1/administrator"
}
```


### Error

```
Status: Default, Ontap error response codes | error code | description | |
---------- | ----------- | | 262278     | user is a required field | |
655865     | the specified path cannot be used if the file does not exist
| | 655470     | the specified share does not exist |
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

share

| Name | Type | Description |
|---|---|---|
| name | string | Displays the file or directory effective permission for the mentioned user, only for files and directories contained where the specified path is relative to the root of the specified share. If this parameter is not specified, the SVM root volume is taken as the default. If this parameter is specified, the effective share permission of the user is also displayed. Wildcard query characters are not supported. |
| path | string | Displays the CIFS share path. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm_reference

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage file security permissions and audit policies

### Manage file security permissions and audit policies

#### Overview

Using this API, You can manage NTFS file security and audit policies of file or directory without the need of a client. It works similar to what you could do with a cacls in windows client. It will create an NTFS security descriptor(SD) to which you can add access control entries (ACEs) to the discretionary access control list (DACL) and the system access control list (SACL). Generally, an SD contains following information:

- Security identifiers (SIDs) for the owner and primary group of an object. A security identifier (SID) is a unique value of variable length used to identify a trustee. Each account has a unique SID issued by an authority, such as a Windows domain controller, and is stored in a security database.
- A DACL identifies the trustees that are allowed or denied access to a securable object. When a process tries to access a securable object, the system checks the ACEs in the object's DACL to determine whether to grant access to it.
- A SACL enables administrators to log attempts to access a secured object. Each ACE specifies the types of access attempts by a specified trustee that cause the system to generate a record in the security event log. An ACE in a SACL can generate audit records when an access attempt fails, when it succeeds, or both.
- A set of control bits that qualify the meaning of a SD or its individual members.

Currently, in ONTAP CLI, creating and applying NTFS ACLs is a 5-step process:

- Create an SD.
- Add DACLs and SACLs to the NTFS SD. If you want to audit file and directory events, you must configure auditing on the Vserver, in addition, to adding a SACL to the SD.

- Create a file/directory security policy. This step associates the policy with a SVM.
- Create a policy task. A policy task refers to a single operation to apply to a file (or folder) or to a set of files (or folders). Among other things, the task defines which SD to apply to a path.
- Apply a policy to the associated SVM. If a permission denied error occurs while applying the policy, the policy is skipped for the relevant files (or folders) and the job continues.

This REST API to set the DACL/SACL is similar to the windows GUI. The approach used here has been simplified by combining all steps into a single step. The REST API uses only minimal and mandatory parameters to create access control entries (ACEs), which can be added to the discretionary access control list (DACL) and the system access control list (SACL). Based on information provided, SD is created and applied on the target path. A path within the FlexCache volume is not supported

Beginning with ONTAP 9.10.1, SLAG (Storage-Level Access Guard) ACLs can also be configured through these endpoints. SLAG is designed to be set on a volume or qtree. Storage-level security cannot be revoked from a client, not even by a system (Windows or UNIX) administrator. It is designed to be modified by storage administrators only, which precedes the share/export permission and the Windows ACLs or UNIX mode bits. Similar to configuring file-directory ACLs, configuring SLAG ACLs is also simplified by combining all steps into a single step.

#### Examples

### Creating a new SD

Use this endpoint to apply a fresh set of SACLs and DACLs. A new SD is created based on the input parameters and it replaces the old SD for the given target path:

```
# The API:
POST /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X POST "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-
0050568e8682/%2Fparent?return_timeout=0" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"acls\": [ { \"access\": \"access_allow\", \"advanced_rights\": {
\"append_data\": true, \"delete\": true, \"delete_child\": true,
\"execute_file\": true, \"full_control\": true, \"read_attr\": true,
\"read_data\": true, \"read_ea\": true, \"read_perm\": true,
\"write_attr\": true, \"write_data\": true, \"write_ea\": true,
\"write_owner\": true, \"write_perm\": true }, \"apply_to\": { \"files\":
true, \"sub_folders\": true, \"this_folder\": true }, \"user\":
\"administrator\" } ], \"control_flags\": \"32788\", \"group\": \"S-1-5-
21-2233347455-2266964949-1780268902-69700\", \"ignore_paths\": [
\"/parent/child2\" ], \"owner\": \"S-1-5-21-2233347455-2266964949-
1780268902-69304\", \"propagation_mode\": \"propagate\"}"

# The response:
{
"job": {
  "uuid": "3015c294-5bbc-11eb-9c4e-0050568e8682",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3015c294-5bbc-11eb-9c4e-0050568e8682"
    }
  }
}
}
```

### Configuring a new set of SLAG DACLs and SACLs

Use this endpoint to apply a fresh set of SLAG DACLs and SACLs. A new SD is created based on the input parameters and it replaces the old SLAG permissions for the given target path:

```
# The API:
POST /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/file-
security/permissions/9f738ac5-c502-11eb-b82c-
0050568e5902/%2Ftest_vol?return_timeout=0" -H  "accept: application/json"
-H  "Content-Type: application/json" -d "{  \"access_control\": \"slag\",
\"acls\": [    {      \"access\": \"access_allow\",
\"advanced_rights\": {        \"append_data\": true,        \"delete\":
true,        \"delete_child\": true,        \"execute_file\": true,
\"full_control\": true,        \"read_attr\": true,        \"read_data\":
true,        \"read_ea\": true,        \"read_perm\": true,
\"write_attr\": true,        \"write_data\": true,        \"write_ea\":
true,        \"write_owner\": true,        \"write_perm\": true      },
\"apply_to\": {        \"files\": true,        \"sub_folders\": true,
\"this_folder\": true      },      \"user\": \"user1\"    },{
\"access\": \"audit_success\",      \"advanced_rights\": {
\"append_data\": true,        \"delete\": true,        \"delete_child\":
true,        \"execute_file\": true,        \"full_control\": true,
\"read_attr\": true,        \"read_data\": true,        \"read_ea\": true,
\"read_perm\": true,        \"write_attr\": true,        \"write_data\":
true,        \"write_ea\": true,        \"write_owner\": true,
\"write_perm\": true      },      \"apply_to\": {        \"files\": true,
\"sub_folders\": true,        \"this_folder\": true      },      \"user\":
\"user2\"    }  ]}"

# The response:
{
"job": {
  "uuid": "9938d743-d566-11eb-ad60-0050568e5902",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9938d743-d566-11eb-ad60-0050568e5902"
    }
  }
}
}
```

### Retrieving file permissions

Use this endpoint to retrieve all the security and auditing information of a directory or file:

```
# The API:
GET /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X GET "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-0050568e8682/%2Fparent" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>"

# The response:
{
"svm": {
  "uuid": "9479099d-5b9f-11eb-9c4e-0050568e8682",
  "name": "vs1"
},
"path": "/parent",
"owner": "BUILTIN\\Administrators",
"group": "BUILTIN\\Administrators",
"control_flags": "0x8014",
"acls": [
  {
    "user": "BUILTIN\\Administrators",
    "access": "access_allow",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
```

```
  },
  {
    "user": "BUILTIN\\Users",
    "access": "access_allow",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
  },
  {
    "user": "CREATOR OWNER",
    "access": "access_allow",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
```

```
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
  },
  {
    "user": "Everyone",
    "access": "access_allow",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
  },
  {
    "user": "NT AUTHORITY\\SYSTEM",
    "access": "access_allow",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
```

```
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
  },
  {
    "user": "user1",
    "access": "access_allow",
    "apply_to": {
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "slag"
  },
  {
    "user": "user1",
    "access": "access_allow",
```

```
    "apply_to": {
      "files": true,
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "slag"
  },
  {
    "user": "user2",
    "access": "audit_success",
    "apply_to": {
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
```

```
      "access_control": "slag"
    },
    {
      "user": "user2",
      "access": "audit_success",
      "apply_to": {
        "files": true,
      },
      "advanced_rights": {
        "append_data": true,
        "delete": true,
        "delete_child": true,
        "execute_file": true,
        "full_control": true,
        "read_attr": true,
        "read_data": true,
        "read_ea": true,
        "read_perm": true,
        "write_attr": true,
        "write_data": true,
        "write_ea": true,
        "write_owner": true,
        "synchronize": true,
        "write_perm": true
      },
      "access_control": "slag"
    }
  ],
  "inode": 64,
  "security_style": "mixed",
  "effective_style": "ntfs",
  "dos_attributes": "10",
  "text_dos_attr": "----D---",
  "user_id": "0",
  "group_id": "0",
  "mode_bits": 777,
  "text_mode_bits": "rwxrwxrwx"
  }
```


#### Configuring lookup_names in a GET request

When the value of lookup_names is set to false (the default value is true), the SID aren’t translated into names:

```
# The API:
GET /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X GET "https://10.140.101.39/api/protocols/file-
security/permissions/bea90fd8-dc69-11ef-9cbd-
005056aeb209/%2Fvol01?lookup_names=false&fields=*" -H  "accept:
application/json" -H  "authorization: Basic <CREDENTIALS>"

# The response:
{
"svm": {
  "uuid": "bea90fd8-dc69-11ef-9cbd-005056aeb209",
  "name": "vs0"
},
"path": "/vol01",
"owner": "S-1-1-0",
"group": "S-1-1-0",
"control_flags": "0x8014",
"acls": [
  {
    "user": "S-1-5-21-1328775103-872486983-3699758552-1002",
    "access": "audit_failure",
    "apply_to": {
      "files": true,
      "sub_folders": true,
      "this_folder": true
    },
    "advanced_rights": {
      "append_data": true,
      "delete": true,
      "delete_child": true,
      "execute_file": true,
      "full_control": true,
      "read_attr": true,
      "read_data": true,
      "read_ea": true,
      "read_perm": true,
      "write_attr": true,
      "write_data": true,
      "write_ea": true,
      "write_owner": true,
      "synchronize": true,
      "write_perm": true
    },
    "access_control": "file_directory"
```

```
    }
  ],
  "inode": 64,
  "security_style": "mixed",
  "effective_style": "ntfs",
  "dos_attributes": "10",
  "text_dos_attr": "----D---",
  "user_id": "0",
  "group_id": "0",
  "mode_bits": 777,
  "text_mode_bits": "rwxrwxrwx"
  }
```


### Updating SD-specific information

Use this end point to update the following information:

- Primary owner of the file/directory.
- Primary group of the file/directory.
- Control flags associated with with SD of the file/directory.

```
# The API:
PATCH /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X PATCH "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-
0050568e8682/%2Fparent?return_timeout=0" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"control_flags\": \"32788\", \"group\": \"everyone\", \"owner\":
\"user1\"}"

# The Response:
{
"job": {
  "uuid": "6f89e612-5bbd-11eb-9c4e-0050568e8682",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6f89e612-5bbd-11eb-9c4e-0050568e8682"
    }
  }
}
}
```

#### Removing all SLAG ACLs

Use this end point to remove all SLAG ACLs.

```
# The API:
DELETE /protocols/file-security/permissions/{svm.uuid}/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/file-
security/permissions/713f569f-d4bc-11eb-b24a-
005056ac6ce1/%2Ftest_vol?access_control=slag"
```

#### Adding a single file-directory DACL/SACL ACE

Use this endpoint to add a single SACL/DACL ACE for a new user or for an existing user with a different access type (allow or deny). The given ACE is merged with an existing SACL/DACL and based on the type of “propagation-mode”, it is reflected to the child object:

```
# The API:
POST /protocols/file-security/permissions/{svm.uuid}/{path}/acl

# The call:
curl -X POST "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-
0050568e8682/%2Fparent/acl?return_timeout=0&return_records=false" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{ \"access\": \"access_allow\",
\"apply_to\": { \"files\": true, \"sub_folders\": true, \"this_folder\":
true }, \"ignore_paths\": [ \"/parent/child2\" ], \"propagation_mode\":
\"propagate\", \"rights\": \"read\", \"user\": \"example_user\"}"

# The Response:
{
"job": {
  "uuid": "26185a2f-5bbe-11eb-9c4e-0050568e8682",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/26185a2f-5bbe-11eb-9c4e-0050568e8682"
    }
  }
}
}
```

### Adding a single SLAG DACL/SACL ACE

Use this endpoint to add a single SLAG SACL/DACL ACE to an existing set of ACLs for a user or for an existing user with a different access type (allow or deny).

```
# The API:
POST /protocols/file-security/permissions/{svm.uuid}/{path}/acl

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/file-
security/permissions/713f569f-d4bc-11eb-b24a-
005056ac6ce1/%2Ftest_vol/acl?return_timeout=0&return_records=false" -H
"accept: application/json" -H  "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{  \"access\": \"access_allow\",
\"access_control\": \"slag\",  \"advanced_rights\": {    \"append_data\":
true,    \"delete\": true,    \"delete_child\": true,    \"execute_file\":
true,    \"full_control\": true,    \"read_attr\": true,    \"read_data\":
true,    \"read_ea\": true,    \"read_perm\": true,    \"write_attr\":
true,    \"write_data\": true,    \"write_ea\": true,    \"write_owner\":
true,    \"write_perm\": true  },  \"apply_to\": {    \"files\": true,
\"sub_folders\": true,    \"this_folder\": true  },  \"user\": \"user1\"}"

# The Response:
{
"job": {
  "uuid": "7fa5f53f-d570-11eb-b24a-005056ac6ce1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/7fa5f53f-d570-11eb-b24a-005056ac6ce1"
    }
  }
}
}
```

#### Updating existing SACL/DACL ACE

Use this endpoint to update the rights/advanced rights for an existing user, for a specified path. You cannot update the access type using this end point. Based on the type of “propagation-mode”, it is reflected to the child object:

```
# The API:
PATCH /protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}
The Call:
curl -X PATCH "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-
0050568e8682/%2Fparent/acl/himanshu?return_timeout=0" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>" -H "Content-
Type: application/json" -d "{ \"access\": \"access_allow\",
\"advanced_rights\": { \"append_data\": true, \"delete\": true,
\"delete_child\": true, \"execute_file\": true, \"full_control\": true,
\"read_attr\": false, \"read_data\": false, \"read_ea\": false,
\"read_perm\": false, \"write_attr\": true, \"write_data\": true,
\"write_ea\": true, \"write_owner\": true, \"write_perm\": true },
\"apply_to\": { \"files\": true, \"sub_folders\": true, \"this_folder\":
true }, \"ignore_paths\": [ \"/parent/child2\" ], \"propagation_mode\":
\"propagate\"}"
The Response:
{
"job": {
  "uuid": "72067401-5bbf-11eb-9c4e-0050568e8682",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/72067401-5bbf-11eb-9c4e-0050568e8682"
    }
  }
}
}
```

### Updating an existing SLG SACL/DACL ACE

Use this endpoint to update the SLAG rights/advanced rights for an existing user, for a specified path. You cannot update the access type using this end point.

```
# The API:
PATCH /protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}
The Call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/file-
security/permissions/713f569f-d4bc-11eb-b24a-
005056ac6ce1/%2Ftest_vol/acl/user1?return_records=false&return_timeout=0"
-H  "accept: application/json" -H  "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{  \"access\": \"access_allow\",
\"access_control\": \"slag\",  \"apply_to\": {    \"files\": true,
\"sub_folders\": true,    \"this_folder\": true  },  \"rights\":
\"read\"}"
The Response:
{
"job": {
  "uuid": "3d21abcd-d571-11eb-b24a-005056ac6ce1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3d21abcd-d571-11eb-b24a-005056ac6ce1"
    }
  }
}
}
```

#### Deleting an existing SACL/DACL ACE

Use this endpoint to delete any of the existing rights/advanced_rights for a user. Based on the type of “propagation-mode”, it is reflected to the child object:

```
# The API:
DELETE /protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}

# The call:
curl -X DELETE "https://10.140.101.39/api/protocols/file-
security/permissions/9479099d-5b9f-11eb-9c4e-
0050568e8682/%2Fparent/acl/himanshu?return_timeout=0" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>" -H "Content-
Type: application/json" -d "{ \"access\": \"access_allow\", \"apply_to\":
{ \"files\": true, \"sub_folders\": true, \"this_folder\": true },
\"ignore_paths\": [ \"/parent/child2\" ], \"propagation_mode\":
\"propagate\"}"

# The response:
{
"job": {
  "uuid": "e5683b61-5bbf-11eb-9c4e-0050568e8682",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e5683b61-5bbf-11eb-9c4e-0050568e8682"
    }
  }
}
}
```

### Deleting an existing SLAG SACL/DACL ACE

Use this endpoint to delete any SLAG ACE for a user.

```
# The API:
DELETE /protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/file-
security/permissions/713f569f-d4bc-11eb-b24a-
005056ac6ce1/%2Ftest_vol/acl/user1?return_records=false&return_timeout=0"
-H  "accept: application/json" -H  "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{  \"access\": \"access_allow\",
\"access_control\": \"slag\",  \"apply_to\": {    \"files\": true,
\"sub_folders\": true,    \"this_folder\": true  }}"

# The response:
{
"job": {
  "uuid": "10c29534-d572-11eb-b24a-005056ac6ce1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/10c29534-d572-11eb-b24a-005056ac6ce1"
    }
  }
}
}
```

---

### Remove all SLAG ACLs for a path

DELETE `/protocols/file-security/permissions/{svm.uuid}/{path}`

**Introduced In:** 9.10

Remove all SLAG ACLs for specified path. Bulk deletion is supported only for SLAG
You must keep the following points in mind while using these endpoints:

- Do not pass additional arguments that are not required.

#### Related ONTAP Commands

- `vserver security file-directory remove-slag`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | target path |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| access_control | string | query | False | Remove all SLAG ACLs. Currently bulk deletion of file-directory ACLs is not supported.<br>• enum: ["slag"] |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 10485811 | Access is a required field. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve file permissions

GET `/protocols/file-security/permissions/{svm.uuid}/{path}`

**Introduced In:** 9.9

Retrieves file permissions.

#### Related ONTAP commands

- `vserver security file-directory show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | target path |
| lookup_names | boolean | query | False | Translate security identifiers into names.<br>• Introduced in: 9.17 |
| fields | array[string] | query | False | Specify the fields to return. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| acls | array[acl] | A discretionary access security list (DACL) identifies the trustees that are allowed or denied access to a securable object. When a process tries to access a securable object, the system checks the access control entries (ACEs) in the object's DACL to determine whether to grant access to it. |
| control_flags | string | Specifies the control flags in the SD. It is a Hexadecimal Value. |

| Name | Type | Description |
| --- | --- | --- |
| dos_attributes | string | Specifies the file attributes on this file or directory. |
| effective_style | string | Specifies the effective style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| group | string | Specifies the owner's primary group.<br>You can specify the owner group using either a group name or SID. |
| group_id | string | Specifies group ID on this file or directory. |
| inode | integer | Specifies the File Inode number. |
| mode_bits | integer | Specifies the mode bits on this file or directory. |
| owner | string | Specifies the owner of the SD.<br>You can specify the owner using either a user name or security identifier (SID).<br>The owner of the SD can modify the permissions on the<br>file (or folder) or files (or folders) to which the SD<br>is applied and can give other users the right to take ownership<br>of the object or objects to which the SD is applied. |
| security_style | string | Specifies the security style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |

| Name | Type | Description |
|---|---|---|
| text_dos_attr | string | Specifies the textual format of file attributes on this file or directory. |
| text_mode_bits | string | Specifies the textual format of mode bits on this file or directory. |
| user_id | string | Specifies user ID of this file or directory. |

#### Example response


```
{
  "acls": [
    {
      "access": "access_allow",
      "access_control": "file_directory",
      "inherited": 1,
      "rights": "full_control",
      "user": "S-1-5-21-2233347455-2266964949-1780268902-69304"
    }
  ],
  "control_flags": "8014",
  "dos_attributes": "10",
  "effective_style": "mixed",
  "group": "S-1-5-21-2233347455-2266964949-1780268902-69700",
  "group_id": "2",
  "inode": 64,
  "mode_bits": 777,
  "owner": "S-1-5-21-2233347455-2266964949-1780268902-69304",
  "security_style": "ntfs",
  "text_dos_attr": "---A----",
  "text_mode_bits": "rwxrwxrwx",
  "user_id": "10"
}
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 655865 | The specified file or directory does not exist. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

advanced_rights

Specifies the advanced access right controlled by the ACE for the account specified.
You can specify more than one "advanced-rights" value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| append_data | boolean | Append DAta |
| delete | boolean | Delete |
| delete_child | boolean | Delete Child |
| execute_file | boolean | Execute File |
| full_control | boolean | Full Control |
| read_attr | boolean | Read Attributes |
| read_data | boolean | Read Data |
| read_ea | boolean | Read Extended Attributes |
| read_perm | boolean | Read Permissions |
| synchronize | boolean | Synchronize |
| write_attr | boolean | Write Attributes |
| write_data | boolean | Write Data |
| write_ea | boolean | Write Extended Attributes |
| write_owner | boolean | Write Owner |
| write_perm | boolean | Write Permission |

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |

| Name | Type | Description |
| --- | --- | --- |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

acl

An ACE is an element in an access control list (ACL).
An ACL can have zero or more ACEs. Each ACE controls or monitors access to an object by a specified trustee.

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• access_allowed_callback - CALLBACK for allowed access<br>• access_denied_callback - CALLBACK for denied access<br>• access_allowed_callback_object - CALLBACK OBJECT for allowed access<br>• access_denied_callback_object - CALLBACK OBJECT for denied access<br>• system_audit_callback - SYSTEM Audit Callback ace<br>• system_audit_callback_object - SYSTEM Audit Callback Object ace<br>• system_resource_attribute - SYSTEM Resource Attribute<br>• system_scoped_policy_id - SYSTEM Scope Policy ID<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access<br>• audit_success_and_failure - SACL for both success and failure access |

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified. You can specify more than one "advanced-rights" value by using a comma-delimited list. |
| inherited | boolean | Indicates whether or not the ACE flag is inherited. |
| rights | string | Specifies the access rights controlled by the ACE for the account specified. The "rights" parameter is mutually exclusive with the "advanced_rights" parameter. If you specify the "rights" parameter, you can specify one of the following "rights" values ("-" or "_" is accepted as the delimiter). |
| user | string | Specifies the account to which the ACE applies. You can specify either name or SID. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the SD information

PATCH `/protocols/file-security/permissions/{svm.uuid}/{path}`

**Introduced In:** 9.9

Updates SD specific Information. For example, owner, group and control-flags. SD specific information of SLAG ACLs is not modifiable.

#### Related ONTAP commands

- `vserver security file-directory ntfs modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | target path |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| control_flags | string | Specifies the control flags in the SD. It is a Hexadecimal Value. |
| dos_attributes | string | Specifies the file attributes on this file or directory. |

| Name | Type | Description |
| --- | --- | --- |
| effective_style | string | Specifies the effective style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| group | string | Specifies the owner's primary group.<br>You can specify the owner group using either a group name or SID. |
| group_id | string | Specifies group ID on this file or directory. |
| inode | integer | Specifies the File Inode number. |
| mode_bits | integer | Specifies the mode bits on this file or directory. |
| owner | string | Specifies the owner of the SD.<br>You can specify the owner using either a user name or security identifier (SID).<br>The owner of the SD can modify the permissions on the<br>file (or folder) or files (or folders) to which the SD<br>is applied and can give other users the right to take ownership<br>of the object or objects to which the SD is applied. |
| security_style | string | Specifies the security style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| text_dos_attr | string | Specifies the textual format of file attributes on this file or directory. |

| Name | Type | Description |
|---|---|---|
| text_mode_bits | string | Specifies the textual format of mode bits on this file or directory. |
| user_id | string | Specifies user ID of this file or directory. |

#### Example request

```
{
  "control_flags": "8014",
  "dos_attributes": "10",
  "effective_style": "mixed",
  "group": "S-1-5-21-2233347455-2266964949-1780268902-69700",
  "group_id": "2",
  "inode": 64,
  "mode_bits": 777,
  "owner": "S-1-5-21-2233347455-2266964949-1780268902-69304",
  "security_style": "ntfs",
  "text_dos_attr": "---A----",
  "text_mode_bits": "rwxrwxrwx",
  "user_id": "10"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |
| 10485814 | The value provided for field control_flags is invalid. |

#### Definitions

**See Definitions**

advanced_rights

Specifies the advanced access right controlled by the ACE for the account specified.
You can specify more than one "advanced-rights" value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| append_data | boolean | Append DAta |
| delete | boolean | Delete |
| delete_child | boolean | Delete Child |
| execute_file | boolean | Execute File |
| full_control | boolean | Full Control |
| read_attr | boolean | Read Attributes |
| read_data | boolean | Read Data |
| read_ea | boolean | Read Extended Attributes |
| read_perm | boolean | Read Permissions |
| synchronize | boolean | Synchronize |
| write_attr | boolean | Write Attributes |
| write_data | boolean | Write Data |
| write_ea | boolean | Write Extended Attributes |
| write_owner | boolean | Write Owner |
| write_perm | boolean | Write Permission |

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |

| Name | Type | Description |
| --- | --- | --- |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

acl

An ACE is an element in an access control list (ACL).
An ACL can have zero or more ACEs. Each ACE controls or monitors access to an object by a specified trustee.

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• access_allowed_callback - CALLBACK for allowed access<br>• access_denied_callback - CALLBACK for denied access<br>• access_allowed_callback_object - CALLBACK OBJECT for allowed access<br>• access_denied_callback_object - CALLBACK OBJECT for denied access<br>• system_audit_callback - SYSTEM Audit Callback ace<br>• system_audit_callback_object - SYSTEM Audit Callback Object ace<br>• system_resource_attribute - SYSTEM Resource Attribute<br>• system_scoped_policy_id - SYSTEM Scope Policy ID<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access<br>• audit_success_and_failure - SACL for both success and failure access |

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified. You can specify more than one "advanced-rights" value by using a comma-delimited list. |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries. You can specify more than one value by using a comma-delimited list. |
| inherited | boolean | Indicates whether or not the ACE flag is inherited. |
| rights | string | Specifies the access rights controlled by the ACE for the account specified. The "rights" parameter is mutually exclusive with the "advanced_rights" parameter. If you specify the "rights" parameter, you can specify one of the following "rights" values ("-" or "_" is accepted as the delimiter). |

file_directory_security

Manages New Technology File System (NTFS) security and NTFS audit policies.

| Name | Type | Description |
| --- | --- | --- |
| control_flags | string | Specifies the control flags in the SD. It is a Hexadecimal Value. |
| dos_attributes | string | Specifies the file attributes on this file or directory. |
| effective_style | string | Specifies the effective style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| group | string | Specifies the owner's primary group.<br>You can specify the owner group using either a group name or SID. |
| group_id | string | Specifies group ID on this file or directory. |
| inode | integer | Specifies the File Inode number. |
| mode_bits | integer | Specifies the mode bits on this file or directory. |
| owner | string | Specifies the owner of the SD. You can specify the owner using either a user name or security identifier (SID).<br>The owner of the SD can modify the permissions on the<br>file (or folder) or files (or folders) to which the SD<br>is applied and can give other users the right to take ownership of the object or objects to which the SD is applied. |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Specifies the security style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| text_dos_attr | string | Specifies the textual format of file attributes on this file or directory. |
| text_mode_bits | string | Specifies the textual format of mode bits on this file or directory. |
| user_id | string | Specifies user ID of this file or directory. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Apply an SD to a path

POST `/protocols/file-security/permissions/{svm.uuid}/{path}`

**Introduced In:** 9.9

Applies an SD to the given path.
You must keep the following points in mind while using these endpoints:

- Either SLAG ACL/s or file-directory ACL/s can be configured in one API call. Both cannot be configured in the same API call.
- SLAG applies to all files and/or directories in a volume hence, inheritance is not required to be propagated.
- Set access_control field to slag while configuring SLAG ACLs.
- Set access_control field to file_directory while configuring file-directory ACLs. By Default access_control field is set to file_directory.
- For SLAG, valid apply_to combinations are "this-folder, sub-folders", "files", "this-folder, sub-folders, files".

#### Related ONTAP commands

- `vserver security file-directory ntfs create`
- `vserver security file-directory ntfs dacl add`
- `vserver security file-directory ntfs sacl add`
- `vserver security file-directory policy create`
- `vserver security file-directory policy task add`
- `vserver security file-directory apply`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | target path |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| acls | array[acl] | A discretionary access security list (DACL) identifies the trustees that are allowed or denied access to a securable object. When a process tries to access a securable object, the system checks the access control entries (ACEs) in the object's DACL to determine whether to grant access to it. |
| control_flags | string | Specifies the control flags in the SD. It is a Hexadecimal Value. |
| dos_attributes | string | Specifies the file attributes on this file or directory. |
| effective_style | string | Specifies the effective style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| group | string | Specifies the owner's primary group. You can specify the owner group using either a group name or SID. |

| Name | Type | Description |
| --- | --- | --- |
| group_id | string | Specifies group ID on this file or directory. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| inode | integer | Specifies the File Inode number. |
| mode_bits | integer | Specifies the mode bits on this file or directory. |
| owner | string | Specifies the owner of the SD.<br>You can specify the owner using either a user name or security identifier (SID).<br>The owner of the SD can modify the permissions on the<br>file (or folder) or files (or folders) to which the SD<br>is applied and can give other users the right to take ownership<br>of the object or objects to which the SD is applied. |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Specifies the security style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| text_dos_attr | string | Specifies the textual format of file attributes on this file or directory. |
| text_mode_bits | string | Specifies the textual format of mode bits on this file or directory. |
| user_id | string | Specifies user ID of this file or directory. |

### Example request

```
{
  "access_control": "file_directory",
  "acls": [
    {
      "access": "access_allow",
      "access_control": "file_directory",
      "inherited": 1,
      "rights": "full_control",
      "user": "S-1-5-21-2233347455-2266964949-1780268902-69304"
    }
  ],
  "control_flags": "8014",
  "dos_attributes": "10",
  "effective_style": "mixed",
  "group": "S-1-5-21-2233347455-2266964949-1780268902-69700",
  "group_id": "2",
  "ignore_paths": [
    "/dir1/dir2/",
    "/parent/dir3"
  ],
  "inode": 64,
  "mode_bits": 777,
  "owner": "S-1-5-21-2233347455-2266964949-1780268902-69304",
  "propagation_mode": "propagate",
  "security_style": "ntfs",
  "text_dos_attr": "---A----",
  "text_mode_bits": "rwxrwxrwx",
  "user_id": "10"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 10485813 | All values corresponding to rights cannot be false. |
| 10485815 | The field "acls.access_control" is not allowed with POST method. |
| 10485810 | User is a required field. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |
| 4849677 | Failed to convert SID to a Windows name. Reason: "SecD Error: object not found". |
| 10485814 | The value provided for field control_flags is invalid. |
| 10485813 | All values cannot be false. |

**See Definitions**

advanced_rights

Specifies the advanced access right controlled by the ACE for the account specified.
You can specify more than one "advanced-rights" value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| append_data | boolean | Append DAta |
| delete | boolean | Delete |
| delete_child | boolean | Delete Child |
| execute_file | boolean | Execute File |
| full_control | boolean | Full Control |
| read_attr | boolean | Read Attributes |
| read_data | boolean | Read Data |
| read_ea | boolean | Read Extended Attributes |
| read_perm | boolean | Read Permissions |
| synchronize | boolean | Synchronize |
| write_attr | boolean | Write Attributes |
| write_data | boolean | Write Data |
| write_ea | boolean | Write Extended Attributes |
| write_owner | boolean | Write Owner |
| write_perm | boolean | Write Permission |

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |

| Name | Type | Description |
| --- | --- | --- |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

acl

An ACE is an element in an access control list (ACL).
An ACL can have zero or more ACEs. Each ACE controls or monitors access to an object by a specified trustee.

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• access_allowed_callback - CALLBACK for allowed access<br>• access_denied_callback - CALLBACK for denied access<br>• access_allowed_callback_object - CALLBACK OBJECT for allowed access<br>• access_denied_callback_object - CALLBACK OBJECT for denied access<br>• system_audit_callback - SYSTEM Audit Callback ace<br>• system_audit_callback_object - SYSTEM Audit Callback Object ace<br>• system_resource_attribute - SYSTEM Resource Attribute<br>• system_scoped_policy_id - SYSTEM Scope Policy ID<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access<br>• audit_success_and_failure - SACL for both success and failure access |

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified. You can specify more than one "advanced-rights" value by using a comma-delimited list. |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries. You can specify more than one value by using a comma-delimited list. |
| inherited | boolean | Indicates whether or not the ACE flag is inherited. |
| rights | string | Specifies the access rights controlled by the ACE for the account specified. The "rights" parameter is mutually exclusive with the "advanced_rights" parameter. If you specify the "rights" parameter, you can specify one of the following "rights" values ("-" or "_" is accepted as the delimiter). |

| Name | Type | Description |
| --- | --- | --- |
| user | string | Specifies the account to which the ACE applies.<br>You can specify either name or SID. |

file_directory_security

Manages New Technology File System (NTFS) security and NTFS audit policies.

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the<br>security descriptors are applied on files and directories at the specified path.<br>The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| acls | array[acl] | A discretionary access security list (DACL) identifies the trustees that are allowed or denied access to a securable object. When a process tries to access a securable<br>object, the system checks the access control entries (ACEs) in the<br>object's DACL to determine whether to grant access to it. |
| control_flags | string | Specifies the control flags in the SD. It is a Hexadecimal Value. |
| dos_attributes | string | Specifies the file attributes on this file or directory. |

| Name | Type | Description |
| --- | --- | --- |
| effective_style | string | Specifies the effective style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| group | string | Specifies the owner's primary group.<br>You can specify the owner group using either a group name or SID. |
| group_id | string | Specifies group ID on this file or directory. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| inode | integer | Specifies the File Inode number. |
| mode_bits | integer | Specifies the mode bits on this file or directory. |
| owner | string | Specifies the owner of the SD.<br>You can specify the owner using either a user name or security identifier (SID).<br>The owner of the SD can modify the permissions on the<br>file (or folder) or files (or folders) to which the SD<br>is applied and can give other users the right to take ownership of the object or objects to which the SD is applied. |

| Name | Type | Description |
| --- | --- | --- |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |
| security_style | string | Specifies the security style of the SD. The following values are supported:<br>• unix - UNIX style<br>• ntfs - NTFS style<br>• mixed - Mixed style<br>• unified - Unified style |
| text_dos_attr | string | Specifies the textual format of file attributes on this file or directory. |
| text_mode_bits | string | Specifies the textual format of mode bits on this file or directory. |
| user_id | string | Specifies user ID of this file or directory. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add a new SACL or DACL ACE

POST `/protocols/file-security/permissions/{svm.uuid}/{path}/acl`

**Introduced In:** 9.9

Adds the new SACL/DACL ACE.
You must keep the following points in mind while using these endpoints:

- SLAG applies to all files and/or directories in a volume hence, inheritance is not required to be propagated.
- Set access_control field to slag while adding SLAG ACE.
- Set access_control field to file_directory while adding file-directory ACE. By Default access_control field is set to file_directory.
- For SLAG, valid apply_to combinations are "this-folder, sub-folders", "files", "this-folder, sub-folders, files".

#### Required properties

- `user` - Name of the user to which the ACE applies.
- `access` - Specifies whether the ACE is for DACL or SACL.

#### Related ONTAP commands

- `vserver security file-directory ntfs dacl add`
- `vserver security file-directory ntfs sacl add`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | path |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |
| access_control | string | Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified.<br>You can specify more than one "advanced-rights" value by using a comma-delimited list. |

| Name | Type | Description |
| --- | --- | --- |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |
| rights | string | Specifies the access rights controlled by the ACE for the account specified.<br>The "rights" parameter is mutually exclusive with the "advanced_rights"<br>parameter. If you specify the "rights" parameter, you can specify one<br>of the following "rights" values ("-" or "_" is accepted as the delimiter). |
| user | string | Specifies the account to which the ACE applies.<br>You can specify either name or SID. |

### Example request

```
{
  "access": "access_allow",
  "access_control": "file_directory",
  "ignore_paths": [
    "/dir1/dir2/",
    "/parent/dir3"
  ],
  "propagation_mode": "string",
  "rights": "full_control",
  "user": "S-1-5-21-2233347455-2266964949-1780268902-69304"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 10485811 | Access is a required field. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |
| 10485810 | User is a required field. |
| 10485813 | All values cannot be false. |

#### Definitions

**See Definitions**

advanced_rights

Specifies the advanced access right controlled by the ACE for the account specified.
You can specify more than one "advanced-rights" value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| append_data | boolean | Append DAta |
| delete | boolean | Delete |
| delete_child | boolean | Delete Child |
| execute_file | boolean | Execute File |
| full_control | boolean | Full Control |
| read_attr | boolean | Read Attributes |
| read_data | boolean | Read Data |
| read_ea | boolean | Read Extended Attributes |
| read_perm | boolean | Read Permissions |
| synchronize | boolean | Synchronize |
| write_attr | boolean | Write Attributes |
| write_data | boolean | Write Data |
| write_ea | boolean | Write Extended Attributes |
| write_owner | boolean | Write Owner |
| write_perm | boolean | Write Permission |

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |

| Name | Type | Description |
| --- | --- | --- |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

file_directory_security_acl

Manages the DACLS or SACLS.

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |
| access_control | string | Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified.<br>You can specify more than one "advanced-rights" value by using a comma-delimited list. |

| Name | Type | Description |
| --- | --- | --- |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |
| rights | string | Specifies the access rights controlled by the ACE for the account specified.<br>The "rights" parameter is mutually exclusive with the "advanced_rights"<br>parameter. If you specify the "rights" parameter, you can specify one<br>of the following "rights" values ("-" or "_" is accepted as the delimiter). |
| user | string | Specifies the account to which the ACE applies.<br>You can specify either name or SID. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete a SACL or DACL ACL**

DELETE `/protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}`

**Introduced In:** 9.9

Deletes the SACL/DACL ACL
You must keep the following points in mind while using these endpoints:

- SLAG applies to all files and/or directories in a volume hence, inheritance is not required to be propagated.
- Set access_control field to slag while deleting SLAG ACE.

- Set access_control field to file_directory while deleting file-directory ACE. By Default access_control field is set to file_directory.
- For SLAG, valid apply_to combinations are "this-folder, sub-folders", "files", "this-folder, sub-folders, files".

#### Required properties

- `access` - Specifies whether the ACE is for DACL or SACL.
- `user` - Name of the user to which the ACE applies.

#### Related ONTAP commands

- `vserver security file-directory ntfs dacl remove`
- `vserver security file-directory ntfs sacl remove`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | path |
| user | string | path | True | User Name |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path.<br>The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |

| Name | Type | Description |
| --- | --- | --- |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |

**Example request**

```
{
  "access": "access_allow",
  "access_control": "file_directory",
  "ignore_paths": [
    "/dir1/dir2/",
    "/parent/dir3"
  ],
  "propagation_mode": "string"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 10485811 | Access is a required field. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |
| 10485813 | All values cannot be false. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

acl_delete

Manages the DACLS or SACLS.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |

| Name | Type | Description |
| --- | --- | --- |
| access_control | string | An Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path.<br>The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update SACLs or DACLs

PATCH `/protocols/file-security/permissions/{svm.uuid}/{path}/acl/{user}`

**Introduced In:** 9.9

Updates the SACLs/DACLs
You must keep the following points in mind while using these endpoints:

- SLAG applies to all files and/or directories in a volume hence, inheritance is not required to be propagated.
- Set access_control field to slag while updating SLAG ACE.
- Set access_control field to file_directory while updating file-directory ACE. By Default access_control field is set to file_directory.
- For SLAG, valid apply_to combinations are "this-folder, sub-folders", "files", "this-folder, sub-folders, files".

#### Required properties

- `access` - Specifies whether the ACE is for DACL or SACL.
- `user` - Name of the user to which the ACE applies.

#### Related ONTAP commands

- `vserver security file-directory ntfs dacl modify`
- `vserver security file-directory ntfs sacl modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | path | True | path |
| user | string | path | True | User Name |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |
| access_control | string | Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified.<br>You can specify more than one "advanced-rights" value by using a comma-delimited list. |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |

| Name | Type | Description |
| --- | --- | --- |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |
| rights | string | Specifies the access rights controlled by the ACE for the account specified.<br>The "rights" parameter is mutually exclusive with the "advanced_rights"<br>parameter. If you specify the "rights" parameter, you can specify one<br>of the following "rights" values ("-" or "_" is accepted as the delimiter). |

**Example request**

```
{
  "access": "access_allow",
  "access_control": "file_directory",
  "ignore_paths": [
    "/dir1/dir2/",
    "/parent/dir3"
  ],
  "propagation_mode": "string",
  "rights": "full_control"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 655865 | The specified file or directory does not exist. |
| 10485811 | Access is a required field. |
| 1260882 | Specified SVM not found. |
| 6691623 | User is not authorized. |
| 4849676 | The specified Windows user or group does not exist. |
| 10485813 | All values cannot be false. |

#### Definitions

**See Definitions**

advanced_rights

Specifies the advanced access right controlled by the ACE for the account specified.
You can specify more than one "advanced-rights" value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| append_data | boolean | Append DAta |
| delete | boolean | Delete |
| delete_child | boolean | Delete Child |
| execute_file | boolean | Execute File |
| full_control | boolean | Full Control |
| read_attr | boolean | Read Attributes |
| read_data | boolean | Read Data |
| read_ea | boolean | Read Extended Attributes |
| read_perm | boolean | Read Permissions |
| synchronize | boolean | Synchronize |
| write_attr | boolean | Write Attributes |
| write_data | boolean | Write Data |
| write_ea | boolean | Write Extended Attributes |
| write_owner | boolean | Write Owner |
| write_perm | boolean | Write Permission |

apply_to

Specifies where to apply the DACL or SACL entries.
You can specify more than one value by using a comma-delimited list.

| Name | Type | Description |
|---|---|---|
| files | boolean | Apply to Files |

| Name | Type | Description |
| --- | --- | --- |
| sub_folders | boolean | Apply to all sub-folders |
| this_folder | boolean | Apply only to this folder |

file_directory_security_acl

Manages the DACLS or SACLS.

| Name | Type | Description |
| --- | --- | --- |
| access | string | Specifies whether the ACL is for DACL or SACL. It is a required field.<br>The available values are:<br>• access_allow - DACL for allow access<br>• access_deny - DACL for deny access<br>• audit_success - SACL for success access<br>• audit_failure - SACL for failure access |
| access_control | string | Access Control Level specifies the access control of the task to be applied. Valid values are "file-directory" or "Storage-Level Access Guard (SLAG)". SLAG is used to apply the specified security descriptors with the task for the volume or qtree. Otherwise, the security descriptors are applied on files and directories at the specified path. The value SLAG is not supported on FlexGroups volumes. The default value is "file-directory" ('-' and '_' are interchangeable). |
| advanced_rights | advanced_rights | Specifies the advanced access right controlled by the ACE for the account specified.<br>You can specify more than one "advanced-rights" value by using a comma-delimited list. |

| Name | Type | Description |
| --- | --- | --- |
| apply_to | apply_to | Specifies where to apply the DACL or SACL entries.<br>You can specify more than one value by using a comma-delimited list. |
| ignore_paths | array[string] | Specifies that permissions on this file or directory cannot be replaced. |
| propagation_mode | string | Specifies how to propagate security settings to child subfolders and files.<br>This setting determines how child files/folders contained within a parent<br>folder inherit access control and audit information from the parent folder.<br>The available values are:<br>• propagate - propagate inheritable permissions to all subfolders and files<br>• ignore - ignore inheritable permissions<br>• replace - replace existing permissions on all subfolders and files with inheritable permissions |
| rights | string | Specifies the access rights controlled by the ACE for the account specified.<br>The "rights" parameter is mutually exclusive with the "advanced_rights"<br>parameter. If you specify the "rights" parameter, you can specify one<br>of the following "rights" values ("-" or "_" is accepted as the delimiter). |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FPolicy configuration

### Manage FPolicy configuration

#### Overview

FPolicy is an infrastructure component of ONTAP that enables partner applications to connect to ONTAP in order to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access. This enables partners to set restrictions on files that are created or accessed on the storage system. FPolicy also allows you to create file policies that specify file operation permissions according to file type. For example, you can restrict certain file types, such as .jpeg and .mp3 files, from being stored on the storage system. FPolicy can monitor file access from CIFS and NFS clients.

As part of FPolicy configuration, you can specify an FPolicy engine which defines the external FPolicy server, FPolicy events, which defines the protocol and file operations to monitor and the FPolicy policy that acts as a container for the FPolicy engine and FPolicy events. It provides a way for policy management functions, such as policy enabling and disabling.

#### Examples

### Creating an FPolicy configuration

To create an FPolicy for an SVM use the following API. Note that the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /protocols/fpolicy/

#The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"engines\": [ { \"name\": \"engine1\", \"port\": 9876,
\"primary_servers\": [ \"10.132.145.22\", \"10.140.101.109\" ],
\"secondary_servers\": [ \"10.132.145.20\", \"10.132.145.21\" ], \"type\":
\"synchronous\", \"format\": \"xml\" } ], \"events\": [ {
\"file_operations\": { \"read\": true, \"write\": true }, \"filters\": {
\"monitor_ads\": true }, \"name\": \"event_cifs\", \"protocol\": \"cifs\",
\"volume_monitoring\": true } ], \"policies\": [ { \"engine\": { \"name\":
\"engine1\" }, \"events\": [ { \"name\": \"event_cifs\" } ],
\"mandatory\": true, \"name\": \"pol0\", \"priority\": 1, \"scope\": {
\"include_volumes\": [ \"vol1\" ] } } ], \"persistent_stores\": [ {
\"name\": \"ps1\", \"volume\": \"psvol\" ,\"size\": 1073741824,
\"autosize_mode\": \"off\"} ], \"svm\": { \"name\": \"vs1\", \"uuid\":
\"b34f5e3d-01d0-11e9-8f63-0050568ea311\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",
      "name": "vs1"
    },
    "engines": [
      {
        "name": "engine1",
        "primary_servers": [
          "10.132.145.22",
          "10.140.101.109"
        ],
        "secondary_servers": [
          "10.132.145.20",
          "10.132.145.21"
        ],
```

```
        "type": "synchronous",
        "port": 9876,
        "format": "xml"
      }
    ],
      "events": [
      {
        "name": "event_cifs",
        "protocol": "cifs",
        "volume_monitoring": true,
        "file_operations": {
          "read": true,
          "write": true
        },
        "filters": {
          "monitor_ads": true
        }
      }
    ],
    "policies": [
      {
        "name": "pol0",
        "priority": 1,
        "events": [
          {
            "name": "event_cifs"
          }
        ],
        "engine": {
          "name": "engine1"
        },
        "scope": {
          "include_volumes": [
            "vol1"
          ]
        },
        "mandatory": true
      }
    ],
    "persistent_stores": [
      {
        "name": "ps1",
        "volume": "psvol",
        "size": 1073741824,
        "autosize_mode": "off"
      }
```

```
    ]
  }
]
}
```

### Retrieving the FPolicy configuration for all the SVMs in the cluster

```
# The API:
GET /protocols/fpolicy

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/fpolicy?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",
      "name": "vs1"
    },
    "engines": [
      {
        "name": "engine1",
        "primary_servers": [
          "10.132.145.22",
          "10.140.101.109"
        ],
        "secondary_servers": [
          "10.132.145.20",
          "10.132.145.21"
        ],
        "type": "synchronous",
        "port": 9876,
        "format": "xml"
      }
    ],
    "events": [
      {
        "name": "event_cifs",
```

```
        "protocol": "cifs",
        "volume_monitoring": true,
        "file_operations": {
          "close": false,
          "create": false,
          "create_dir": false,
          "delete": false,
          "delete_dir": false,
          "getattr": false,
          "link": false,
          "lookup": false,
          "open": false,
          "read": true,
          "write": true,
          "rename": false,
          "rename_dir": false,
          "setattr": false,
          "symlink": false
        },
        "filters": {
          "monitor_ads": true,
          "close_with_modification": false,
          "close_without_modification": false,
          "close_with_read": false,
          "first_read": false,
          "first_write": false,
          "offline_bit": false,
          "open_with_delete_intent": false,
          "open_with_write_intent": false,
          "write_with_size_change": false,
          "setattr_with_owner_change": false,
          "setattr_with_group_change": false,
          "setattr_with_sacl_change": false,
          "setattr_with_dacl_change": false,
          "setattr_with_modify_time_change": false,
          "setattr_with_access_time_change": false,
          "setattr_with_creation_time_change": false,
          "setattr_with_mode_change": false,
          "setattr_with_size_change": false,
          "setattr_with_allocation_size_change": false,
          "exclude_directory": false
        }
      }
    ],
    "policies": [
      {
```

```
        "name": "pol0",
        "enabled": true,
        "priority": 1,
        "events": [
          {
            "name": "event_cifs"
          }
        ],
        "engine": {
          "name": "engine1"
        },
        "scope": {
          "include_volumes": [
            "vol1"
          ]
        },
        "mandatory": true,
        "passthrough_read": false,
        "allow_privileged_access": false,
        "persistent_store": "ps1"
      }
    ],
    "persistent_stores": [
      {
        "name": "ps1",
        "volume": "psvol",
        "size": 1073741824,
        "autosize_mode": "off"
      }
    ]
  }
],
"num_records": 1
}
```


#### Retrieving an FPolicy configuration for a particular SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/b34f5e3d-01d0-11e9-
```

```
8f63-0050568ea311?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"svm": {
  "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",
  "name": "vs1"
},
"engines": [
  {
    "name": "engine1",
    "primary_servers": [
      "10.132.145.22",
      "10.140.101.109"
    ],
    "secondary_servers": [
      "10.132.145.20",
      "10.132.145.21"
    ],
    "type": "synchronous",
    "port": 9876,
    "format": "xml"
  }
],
"events": [
  {
    "name": "event_cifs",
    "protocol": "cifs",
    "volume_monitoring": true,
    "file_operations": {
      "close": false,
      "create": false,
      "create_dir": false,
      "delete": false,
      "delete_dir": false,
      "getattr": false,
      "link": false,
      "lookup": false,
      "open": false,
      "read": true,
      "write": true,
      "rename": false,
      "rename_dir": false,
      "setattr": false,
      "symlink": false
```

```
    },
    "filters": {
      "monitor_ads": true,
      "close_with_modification": false,
      "close_without_modification": false,
      "close_with_read": false,
      "first_read": false,
      "first_write": false,
      "offline_bit": false,
      "open_with_delete_intent": false,
      "open_with_write_intent": false,
      "write_with_size_change": false,
      "setattr_with_owner_change": false,
      "setattr_with_group_change": false,
      "setattr_with_sacl_change": false,
      "setattr_with_dacl_change": false,
      "setattr_with_modify_time_change": false,
      "setattr_with_access_time_change": false,
      "setattr_with_creation_time_change": false,
      "setattr_with_mode_change": false,
      "setattr_with_size_change": false,
      "setattr_with_allocation_size_change": false,
      "exclude_directory": false
    }
  }
],
"policies": [
  {
    "name": "pol0",
    "enabled": true,
    "priority": 1,
    "events": [
      {
        "name": "event_cifs"
      }
    ],
    "engine": {
      "name": "engine1"
    },
    "scope": {
      "include_volumes": [
        "vol1"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
```

```
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  }
],
"persistent_stores": [
  {
    "name": "ps1",
    "volume": "psvol",
    "size": 1073741824,
    "autosize_mode": "off"
  }
]
}
```

### Deleting an FPolicy configuration for a particular SVM

```
# The API:
DELETE /protocols/fpolicy/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/b34f5e3d-01d0-
11e9-8f63-0050568ea311" -H "accept: application/json"
```

### Retrieve an FPolicy configuration

GET `/protocols/fpolicy`

**Introduced In:** 9.6

Retrieves an FPolicy configuration.

#### Related ONTAP commands

- `fpolicy show`
- `fpolicy policy show`
- `fpolicy policy scope show`
- `fpolicy policy event show`
- `fpolicy policy external-engine show`
- `fpolicy persistent-store show`

**Learn more**

- DOC /protocols/fpolicy

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policies.name | string | query | False | Filter by policies.name |
| policies.priority | integer | query | False | Filter by policies.priority<br>• Max value: 10<br>• Min value: 1 |
| policies.privileged_user | string | query | False | Filter by policies.privileged_user<br>• Introduced in: 9.11 |
| policies.allow_privileged_access | boolean | query | False | Filter by policies.allow_privileged_access<br>• Introduced in: 9.13 |
| policies.enabled | boolean | query | False | Filter by policies.enabled |
| policies.engine.name | string | query | False | Filter by policies.engine.name |
| policies.events.name | string | query | False | Filter by policies.events.name |
| policies.persistent_store | string | query | False | Filter by policies.persistent_store<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policies.scope.include_shares | string | query | False | Filter by policies.scope.include_shares |
| policies.scope.exclude_shares | string | query | False | Filter by policies.scope.exclude_shares |
| policies.scope.exclude_extension | string | query | False | Filter by policies.scope.exclude_extension |
| policies.scope.object_monitoring_with_no_extension | boolean | query | False | Filter by policies.scope.object_monitoring_with_no_extension<br>• Introduced in: 9.11 |
| policies.scope.exclude_volumes | string | query | False | Filter by policies.scope.exclude_volumes |
| policies.scope.exclude_export_policies | string | query | False | Filter by policies.scope.exclude_export_policies |
| policies.scope.include_export_policies | string | query | False | Filter by policies.scope.include_export_policies |
| policies.scope.include_extension | string | query | False | Filter by policies.scope.include_extension |
| policies.scope.include_volumes | string | query | False | Filter by policies.scope.include_volumes |
| policies.scope.check_extensions_on_directories | boolean | query | False | Filter by policies.scope.check_extensions_on_directories<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policies.passthrough_read | boolean | query | False | Filter by policies.passthrough_read<br>• Introduced in: 9.11 |
| policies.mandatory | boolean | query | False | Filter by policies.mandatory |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| events.protocol | string | query | False | Filter by events.protocol |
| events.file_operations.getattr | boolean | query | False | Filter by events.file_operations.getattr |
| events.file_operations.read | boolean | query | False | Filter by events.file_operations.read |
| events.file_operations.write | boolean | query | False | Filter by events.file_operations.write |
| events.file_operations.close | boolean | query | False | Filter by events.file_operations.close |
| events.file_operations.setattr | boolean | query | False | Filter by events.file_operations.setattr |
| events.file_operations.create | boolean | query | False | Filter by events.file_operations.create |
| events.file_operations.create_dir | boolean | query | False | Filter by events.file_operations.create_dir |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| events.file_operations.lookup | boolean | query | False | Filter by events.file_operations.lookup |
| events.file_operations.delete_dir | boolean | query | False | Filter by events.file_operations.delete_dir |
| events.file_operations.delete | boolean | query | False | Filter by events.file_operations.delete |
| events.file_operations.access | boolean | query | False | Filter by events.file_operations.access<br>• Introduced in: 9.13 |
| events.file_operations.open | boolean | query | False | Filter by events.file_operations.open |
| events.file_operations.rename_dir | boolean | query | False | Filter by events.file_operations.rename_dir |
| events.file_operations.symlink | boolean | query | False | Filter by events.file_operations.symlink |
| events.file_operations.link | boolean | query | False | Filter by events.file_operations.link |
| events.file_operations.rename | boolean | query | False | Filter by events.file_operations.rename |
| events.filters.setattr_with_access_time_change | boolean | query | False | Filter by events.filters.setattr_with_access_time_change |
| events.filters.open_with_delete_intent | boolean | query | False | Filter by events.filters.open_with_delete_intent |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| events.filters.close_without_modification | boolean | query | False | Filter by events.filters.close_without_modification |
| events.filters.setattr_with_size_change | boolean | query | False | Filter by events.filters.setattr_with_size_change |
| events.filters.close_with_read | boolean | query | False | Filter by events.filters.close_with_read |
| events.filters.write_with_size_change | boolean | query | False | Filter by events.filters.write_with_size_change |
| events.filters.setattr_with_owner_change | boolean | query | False | Filter by events.filters.setattr_with_owner_change |
| events.filters.exclude_directory | boolean | query | False | Filter by events.filters.exclude_directory |
| events.filters.first_read | boolean | query | False | Filter by events.filters.first_read |
| events.filters.first_write | boolean | query | False | Filter by events.filters.first_write |
| events.filters.setattr_with_modify_time_change | boolean | query | False | Filter by events.filters.setattr_with_modify_time_change |
| events.filters.monitor_ads | boolean | query | False | Filter by events.filters.monitor_ads |
| events.filters.setattr_with_dacl_change | boolean | query | False | Filter by events.filters.setattr_with_dacl_change |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| events.filters.close_with_modification | boolean | query | False | Filter by events.filters.close_with_modification |
| events.filters.setattr_with_creation_time_change | boolean | query | False | Filter by events.filters.setattr_with_creation_time_change |
| events.filters.offline_bit | boolean | query | False | Filter by events.filters.offline_bit |
| events.filters.open_with_write_intent | boolean | query | False | Filter by events.filters.open_with_write_intent |
| events.filters.setattr_with_sacl_change | boolean | query | False | Filter by events.filters.setattr_with_sacl_change |
| events.filters.setattr_with_group_change | boolean | query | False | Filter by events.filters.setattr_with_group_change |
| events.filters.setattr_with_mode_change | boolean | query | False | Filter by events.filters.setattr_with_mode_change |
| events.filters.setattr_with_allocation_size_change | boolean | query | False | Filter by events.filters.setattr_with_allocation_size_change |
| events.name | string | query | False | Filter by events.name |
| events.monitor_fileop_failure | boolean | query | False | Filter by events.monitor_fileop_failure<br>• Introduced in: 9.13 |
| events.volume_monitoring | boolean | query | False | Filter by events.volume_monitoring |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| persistent_stores.volume | string | query | False | Filter by persistent_stores.volume<br><br>• Introduced in: 9.14 |
| persistent_stores.name | string | query | False | Filter by persistent_stores.name<br><br>• Introduced in: 9.14 |
| persistent_stores.size | integer | query | False | Filter by persistent_stores.size<br><br>• Introduced in: 9.15 |
| persistent_stores.autosize_mode | string | query | False | Filter by persistent_stores.autosize_mode<br><br>• Introduced in: 9.15 |
| engines.keep_alive_interval | string | query | False | Filter by engines.keep_alive_interval<br><br>• Introduced in: 9.13 |
| engines.request_cancel_timeout | string | query | False | Filter by engines.request_cancel_timeout<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engines.max_server_requests | integer | query | False | Filter by engines.max_server_requests<br>• Introduced in: 9.11<br>• Max value: 10000<br>• Min value: 1 |
| engines.port | integer | query | False | Filter by engines.port |
| engines.type | string | query | False | Filter by engines.type |
| engines.resiliency.enabled | boolean | query | False | Filter by engines.resiliency.enabled<br>• Introduced in: 9.11 |
| engines.resiliency.retention_duration | string | query | False | Filter by engines.resiliency.retention_duration<br>• Introduced in: 9.11 |
| engines.resiliency.directory_path | string | query | False | Filter by engines.resiliency.directory_path<br>• Introduced in: 9.11 |
| engines.server_progress_timeout | string | query | False | Filter by engines.server_progress_timeout<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engines.ssl_option | string | query | False | Filter by engines.ssl_option<br>• Introduced in: 9.11 |
| engines.status_request_interval | string | query | False | Filter by engines.status_request_interval<br>• Introduced in: 9.11 |
| engines.buffer_size.recv_buffer | integer | query | False | Filter by engines.buffer_size.recv_buffer<br>• Introduced in: 9.11<br>• Max value: 7895160<br>• Min value: 0 |
| engines.buffer_size.send_buffer | integer | query | False | Filter by engines.buffer_size.send_buffer<br>• Introduced in: 9.11<br>• Max value: 7895160<br>• Min value: 0 |
| engines.primary_servers | string | query | False | Filter by engines.primary_servers |
| engines.name | string | query | False | Filter by engines.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engines.max_connection_retries | integer | query | False | Filter by engines.max_connection_retries<br>• Introduced in: 9.17<br>• Max value: 20<br>• Min value: 0 |
| engines.certificate.serial_number | string | query | False | Filter by engines.certificate.serial_number<br>• Introduced in: 9.11 |
| engines.certificate.ca | string | query | False | Filter by engines.certificate.ca<br>• Introduced in: 9.11 |
| engines.certificate.name | string | query | False | Filter by engines.certificate.name<br>• Introduced in: 9.11 |
| engines.session_timeout | string | query | False | Filter by engines.session_timeout<br>• Introduced in: 9.17 |
| engines.secondary_servers | string | query | False | Filter by engines.secondary_servers |
| engines.format | string | query | False | Filter by engines.format<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engines.request_abort_timeout | string | query | False | Filter by engines.request_abort_timeout<br>• Introduced in: 9.11 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[fpolicy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "engines": [
        {
          "certificate": {
            "ca": "TASample1",
            "name": "Sample1-FPolicy-Client",
            "serial_number": "8DDE112A114D1FBC"
          },
          "format": "string",
          "keep_alive_interval": "PT2M",
          "max_connection_retries": 5,
          "max_server_requests": 500,
          "name": "fp_ex_eng",
          "port": 9876,
          "primary_servers": [
            "10.132.145.20",
            "10.140.101.109"
          ],
          "request_abort_timeout": "PT40S",
          "request_cancel_timeout": "PT20S",
          "resiliency": {
            "directory_path": "/dir1",
            "retention_duration": "PT3M"
          },
          "secondary_servers": [
            "10.132.145.20",
            "10.132.145.21"
          ],
```

```
          "server_progress_timeout": "PT1M",
          "session_timeout": "PT10S",
          "ssl_option": "string",
          "status_request_interval": "PT10S",
          "type": "string"
        }
      ],
      "events": [
        {
          "name": "event_cifs",
          "protocol": "string"
        }
      ],
      "persistent_stores": [
        {
          "autosize_mode": "string",
          "name": "ps1",
          "size": "100M",
          "volume": "psvol"
        }
      ],
      "policies": [
        {
          "engine": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "string"
          },
          "events": [
            "event_cifs",
            "event_open"
          ],
          "name": "fp_policy_1",
          "persistent_store": "ps1",
          "priority": 1,
          "privileged_user": "mydomain\\testuser",
          "scope": {
            "exclude_export_policies": [
              "string"
            ],
            "exclude_extension": [
              "string"
            ],
```

```
            "exclude_shares": [
              "string"
            ],
            "exclude_volumes": [
              "vol1",
              "vol_svm1",
              "*"
            ],
            "include_export_policies": [
              "string"
            ],
            "include_extension": [
              "string"
            ],
            "include_shares": [
              "sh1",
              "share_cifs"
            ],
            "include_volumes": [
              "vol1",
              "vol_svm1"
            ]
          }
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

fpolicy_engines

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |

| Name | Type | Description |
| --- | --- | --- |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
| --- | --- | --- |
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |

| Name | Type | Description |
| --- | --- | --- |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |

| Name | Type | Description |
| --- | --- | --- |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

fpolicy_events

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

fpolicy_persistent_stores

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
| --- | --- | --- |
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

scope

| Name | Type | Description |
| --- | --- | --- |
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true, all objects with or without extensions are monitored. Default is false. |

fpolicy_policies

| Name | Type | Description |
| --- | --- | --- |
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes.<br>When this parameter is set to true,<br>FPolicy servers can access files on the cluster using a separate data channel with privileged access. |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period.<br>When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy.<br>Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |

| Name | Type | Description |
| --- | --- | --- |
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster<br>using a separate data channel with privileged access. The input for<br>this field should be in "domain\username" format. |
| scope | scope | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fpolicy

FPolicy is an infrastructure component of ONTAP that enables partner applications connected to your storage systems to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| engines | array[fpolicy_engines] | |
| events | array[fpolicy_events] | |
| persistent_stores | array[fpolicy_persistent_stores] | |
| policies | array[fpolicy_policies] | |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an FPolicy configuration

POST `/protocols/fpolicy`

**Introduced In:** 9.6

Creates an FPolicy configuration.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the FPolicy configuration.

#### Recommended optional properties

- `engines` - External server to which the notifications will be sent.
- `events` - File operations to monitor.
- `policies` - Policy configuration which acts as a container for FPolicy event and FPolicy engine.
- `scope` - Scope of the policy. Can be limited to exports, volumes, shares or file extensions.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `engines.type` - *synchronous*
- `policies.engine` - *native*
- `policies.mandatory` - *true*
- `events.volume_monitoring` - *false*
- `events.file_operations.*` - *false*
- `events.filters.*` - *false*
- `events.monitor_fileop_failure.*` - *false*

#### Related ONTAP commands

- `fpolicy policy event create`
- `fpolicy policy external-engine create`
- `fpolicy policy create`
- `fpolicy policy scope create`
- `fpolicy enable`
- `fpolicy persistent-store create`

#### Learn more

- DOC /protocols/fpolicy

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| engines | array[fpolicy_engines] | |
| events | array[fpolicy_events] | |
| persistent_stores | array[fpolicy_persistent_stores] | |
| policies | array[fpolicy_policies] | |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "engines": [
    {
      "certificate": {
        "ca": "TASample1",
        "name": "Sample1-FPolicy-Client",
        "serial_number": "8DDE112A114D1FBC"
      },
      "format": "string",
      "keep_alive_interval": "PT2M",
      "max_connection_retries": 5,
      "max_server_requests": 500,
      "name": "fp_ex_eng",
      "port": 9876,
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "request_abort_timeout": "PT40S",
      "request_cancel_timeout": "PT20S",
      "resiliency": {
        "directory_path": "/dir1",
        "retention_duration": "PT3M"
      },
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "server_progress_timeout": "PT1M",
      "session_timeout": "PT10S",
      "ssl_option": "string",
      "status_request_interval": "PT10S",
      "type": "string"
    }
  ],
  "events": [
    {
      "name": "event_cifs",
      "protocol": "string"
    }
  ],
  "persistent_stores": [
    {
      "autosize_mode": "string",
```

```
      "name": "ps1",
      "size": "100M",
      "volume": "psvol"
    }
  ],
  "policies": [
    {
      "allow_privileged_access": true,
      "enabled": true,
      "engine": {
        "name": "string"
      },
      "events": [
        "event_cifs",
        "event_open"
      ],
      "name": "fp_policy_1",
      "persistent_store": "ps1",
      "priority": 1,
      "privileged_user": "mydomain\\testuser",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
          "string"
        ],
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [
```

```
            "vol1",
            "vol_svm1"
          ]
        }
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[fpolicy] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "engines": [
        {
          "certificate": {
            "ca": "TASample1",
            "name": "Sample1-FPolicy-Client",
            "serial_number": "8DDE112A114D1FBC"
          },
          "format": "string",
          "keep_alive_interval": "PT2M",
          "max_connection_retries": 5,
          "max_server_requests": 500,
          "name": "fp_ex_eng",
          "port": 9876,
          "primary_servers": [
            "10.132.145.20",
            "10.140.101.109"
          ],
          "request_abort_timeout": "PT40S",
          "request_cancel_timeout": "PT20S",
          "resiliency": {
            "directory_path": "/dir1",
            "retention_duration": "PT3M"
          },
          "secondary_servers": [
            "10.132.145.20",
            "10.132.145.21"
          ],
          "server_progress_timeout": "PT1M",
          "session_timeout": "PT10S",
          "ssl_option": "string",
          "status_request_interval": "PT10S",
          "type": "string"
        }
      ],
      "events": [
        {
          "name": "event_cifs",
          "protocol": "string"
        }
      ],
```

```
    "persistent_stores": [
      {
        "autosize_mode": "string",
        "name": "ps1",
        "size": "100M",
        "volume": "psvol"
      }
    ],
    "policies": [
      {
        "allow_privileged_access": true,
        "enabled": true,
        "engine": {
          "name": "string"
        },
        "events": [
          "event_cifs",
          "event_open"
        ],
        "name": "fp_policy_1",
        "persistent_store": "ps1",
        "priority": 1,
        "privileged_user": "mydomain\\testuser",
        "scope": {
          "exclude_export_policies": [
            "string"
          ],
          "exclude_extension": [
            "string"
          ],
          "exclude_shares": [
            "string"
          ],
          "exclude_volumes": [
            "vol1",
            "vol_svm1",
            "*"
          ],
          "include_export_policies": [
            "string"
          ],
          "include_extension": [
            "string"
          ],
          "include_shares": [
            "sh1",
```

```
                "share_cifs"
              ],
              "include_volumes": [
                "vol1",
                "vol_svm1"
              ]
            }
          }
        ],
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9765032 | The FPolicy engine, FPolicy event or FPolicy policy specified already exists |
| 9765031 | If any of the FPolicy engine, FPolicy event, or FPolicy policy creation fails due to a systematic error or hardware failure, the cause of the failure is detailed in the error message |
| 2621706 | The SVM UUID specified belongs to different SVM |
| 2621462 | The SVM name specified does not exist |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

fpolicy_engines

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |

| Name | Type | Description |
| --- | --- | --- |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
| --- | --- | --- |
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |

| Name | Type | Description |
| --- | --- | --- |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |

| Name | Type | Description |
| --- | --- | --- |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

fpolicy_events

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

fpolicy_persistent_stores

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
| --- | --- | --- |
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

scope

| Name | Type | Description |
| --- | --- | --- |
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true, all objects with or without extensions are monitored. Default is false. |

fpolicy_policies

| Name | Type | Description |
| --- | --- | --- |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |

| Name | Type | Description |
| --- | --- | --- |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers registered for the policy. Passthrough-read is a way to read data for offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fpolicy

FPolicy is an infrastructure component of ONTAP that enables partner applications connected to your storage systems to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access.

| Name | Type | Description |
| --- | --- | --- |
| engines | array[fpolicy_engines] | |
| events | array[fpolicy_events] | |
| persistent_stores | array[fpolicy_persistent_stores] | |
| policies | array[fpolicy_policies] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the FPolicy configuration for an SVM

DELETE `/protocols/fpolicy/{svm.uuid}`

**Introduced In:** 9.6

Deletes the FPolicy configuration for the specified SVM. Before deleting the FPolicy configuration, ensure that all policies belonging to the SVM are disabled.

Important notes:
The volume associated with the FPolicy Persistent Store is not automatically deleted from the FPolicy general endpoint. The associated volume can be removed manually.

#### Related ONTAP commands

- `fpolicy delete`
- `fpolicy policy scope delete`
- `fpolicy policy delete`
- `fpolicy policy event delete`
- `fpolicy policy external-engine delete`
- `fpolicy persistent-store delete`

#### Learn more

- DOC /protocols/fpolicy

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 9765031 | If any of the FPolicy engine, FPolicy event or FPolicy policy deletion fails due to a systemic error or hardware failure, the cause of the failure is detailed in the error message. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the FPolicy configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}`

**Introduced In:** 9.6

Retrieves an FPolicy configuration of an SVM.

#### Related ONTAP commands

- `fpolicy show`
- `fpolicy policy show`
- `fpolicy policy scope show`
- `fpolicy policy event show`
- `fpolicy policy external-engine show`
- `fpolicy persistent-store show`

#### Learn more

- DOC /protocols/fpolicy

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| engines | array[fpolicy_engines] | |
| events | array[fpolicy_events] | |
| persistent_stores | array[fpolicy_persistent_stores] | |
| policies | array[fpolicy_policies] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "engines": [
    {
      "certificate": {
        "ca": "TASample1",
        "name": "Sample1-FPolicy-Client",
        "serial_number": "8DDE112A114D1FBC"
      },
      "format": "string",
      "keep_alive_interval": "PT2M",
      "max_connection_retries": 5,
      "max_server_requests": 500,
      "name": "fp_ex_eng",
      "port": 9876,
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "request_abort_timeout": "PT40S",
      "request_cancel_timeout": "PT20S",
      "resiliency": {
        "directory_path": "/dir1",
        "retention_duration": "PT3M"
      },
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "server_progress_timeout": "PT1M",
      "session_timeout": "PT10S",
      "ssl_option": "string",
      "status_request_interval": "PT10S",
      "type": "string"
    }
  ],
  "events": [
    {
      "name": "event_cifs",
      "protocol": "string"
```

```
    }
  ],
  "persistent_stores": [
    {
      "autosize_mode": "string",
      "name": "ps1",
      "size": "100M",
      "volume": "psvol"
    }
  ],
  "policies": [
    {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "events": [
        "event_cifs",
        "event_open"
      ],
      "name": "fp_policy_1",
      "persistent_store": "ps1",
      "priority": 1,
      "privileged_user": "mydomain\\testuser",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
          "string"
        ],
```

```
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [
          "vol1",
          "vol_svm1"
        ]
      }
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

fpolicy_engines

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |

| Name | Type | Description |
| --- | --- | --- |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.10<br>◦ x-nullable: true |

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
| --- | --- | --- |
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |

| Name | Type | Description |
| --- | --- | --- |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |

| Name | Type | Description |
| --- | --- | --- |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

fpolicy_events

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

fpolicy_persistent_stores

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
| --- | --- | --- |
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

scope

| Name | Type | Description |
| --- | --- | --- |
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true, all objects with or without extensions are monitored. Default is false. |

fpolicy_policies

| Name | Type | Description |
| --- | --- | --- |
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes.<br>When this parameter is set to true,<br>FPolicy servers can access files on the cluster using a separate data channel with privileged access. |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period.<br>When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy.<br>Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |

| Name | Type | Description |
| --- | --- | --- |
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster<br>using a separate data channel with privileged access. The input for<br>this field should be in "domain\username" format. |
| scope | scope | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update FPolicy server configuration

### View and update FPolicy server configuration

#### Overview

This API is used to display and update connection status information for external FPolicy servers.
You must keep the following in mind while using these endpoints:

- If the passthrough_read field is set to true in a GET collection call, only FPolicy passthrough-read connections are returned.
- If the passthrough_read field is not provided or set to false in a GET collection call, only FPolicy server connections are returned.

#### Examples

#### Retrieving the FPolicy server connections for all SVMs in the cluster

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/connections

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/fpolicy/*/connections?passthrough_read=false&fields=*&re
turn_records=true" -H  "accept: application/json"

# The Response:
{
"records": [
  {
    "node": {
      "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
      "name": "user-vsim4"
    },
    "svm": {
      "uuid": "9f738ac5-c502-11eb-b82c-0050568e5902",
```

```
      "name": "vs1"
    },
    "policy": {
      "name": "p1"
    },
    "server": "192.168.137.78",
    "state": "disconnected",
    "update_time": "2021-06-17T16:05:15+05:30",
    "disconnected_reason": {
      "message": "No local lif present to connect to FPolicy server.",
      "code": 9305
    },
    "type": "primary"
  },
  {
    "node": {
      "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
      "name": "user-vsim4"
    },
    "svm": {
      "uuid": "9f738ac5-c502-11eb-b82c-0050568e5902",
      "name": "vs1"
    },
    "policy": {
      "name": "p2"
    },
    "server": "192.168.136.38",
    "state": "disconnected",
    "update_time": "2021-06-17T16:05:15+05:30",
    "disconnected_reason": {
      "message": "No local lif present to connect to FPolicy server.",
      "code": 9305
    },
    "type": "primary"
  },
  {
    "node": {
      "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
      "name": "user-vsim4"
    },
    "svm": {
      "uuid": "b6df362b-c502-11eb-b82c-0050568e5902",
      "name": "vs2"
    },
    "policy": {
      "name": "pol1"
```

```
    },
    "server": "192.168.129.146",
    "state": "disconnected",
    "update_time": "2021-06-17T16:05:15+05:30",
    "disconnected_reason": {
      "message": "No local lif present to connect to FPolicy server.",
      "code": 9305
    },
    "type": "primary"
  }
],
"num_records": 3
}
```

#### Retrieving all FPolicy passthrough read connections for all SVMs in the cluster

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/connections

# The Call:
curl-X GET "https://<mgmt-
ip>/api/protocols/fpolicy/*/connections?passthrough_read=true&fields=*&ret
urn_records=true&return_timeout=15"

# The Response:
{
"records": [
  {
    "node": {
      "uuid": "55693090-c7c8-11eb-a07a-0050568ebc01",
      "name": "user-vsim3",
    },
    "svm": {
      "uuid": "a69e938d-c7ca-11eb-a07a-0050568ebc01",
      "name": "vs2",
    },
    "policy": {
      "name": "pol1",
    },
    "server": "192.168.129.146",
    "session_uuid": "2410d348-c7cb-11eb-a07a-0050568ebc01",
    "state": "connected",
```

```
  },
  {
    "node": {
      "uuid": "55693090-c7c8-11eb-a07a-0050568ebc01",
      "name": "user-vsim3",
    },
    "svm": {
      "uuid": "a69e938d-c7ca-11eb-a07a-0050568ebc01",
      "name": "vs2",
    },
    "policy": {
      "name": "pol2",
    },
    "server": "192.168.129.146",
    "session_uuid": "288f7002-c7cb-11eb-a07a-0050568ebc01",
    "state": "connected",
  }
],
"num_records": 2,
}
```

#### Retrieving the FPolicy server connections for a specific SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/connections

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/9f738ac5-c502-11eb-
b82c-
0050568e5902/connections?passthrough_read=false&fields=*&return_records=tr
ue" -H  "accept: application/json"

# The Response:
{
"records": [
  {
    "node": {
      "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
      "name": "user-vsim4"
    },
    "svm": {
      "uuid": "9f738ac5-c502-11eb-b82c-0050568e5902",
```

```
      "name": "vs1"
    },
    "policy": {
      "name": "p1"
    },
    "server": "192.168.137.78",
    "state": "disconnected",
    "update_time": "2021-06-17T16:05:15+05:30",
    "disconnected_reason": {
      "message": "No local lif present to connect to FPolicy server.",
      "code": 9305
    },
    "type": "primary"
  },
  {
    "node": {
      "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
      "name": "user-vsim4"
    },
    "svm": {
      "uuid": "9f738ac5-c502-11eb-b82c-0050568e5902",
      "name": "vs1"
    },
    "policy": {
      "name": "p2"
    },
    "server": "192.168.136.38",
    "state": "disconnected",
    "update_time": "2021-06-17T16:05:15+05:30",
    "disconnected_reason": {
      "message": "No local lif present to connect to FPolicy server.",
      "code": 9305
    },
    "type": "primary"
  },
],
"num_records": 2
}
```


### Retrieving a specific FPolicy server connection

```
# The API:
GET
/protocols/fpolicy/{svm.uuid}/connections/{node.uuid}/{policy.name}/{serve
r}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/9f738ac5-c502-11eb-
b82c-0050568e5902/connections/8ca36b68-c501-11eb-b82c-
0050568e5902/p1/192.168.137.78" -H  "accept: application/json"

# The Response:
{
"node": {
  "uuid": "8ca36b68-c501-11eb-b82c-0050568e5902",
  "name": "user-vsim4"
},
"svm": {
  "uuid": "9f738ac5-c502-11eb-b82c-0050568e5902",
  "name": "vs1"
},
"policy": {
  "name": "p1"
},
"server": "192.168.137.78",
"state": "disconnected",
"update_time": "2021-06-17T16:05:15+05:30",
"disconnected_reason": {
  "message": "No local lif present to connect to FPolicy server.",
  "code": 9305
},
"type": "primary"
}
```

### Updating the FPolicy server connection

```
# The API:
PATCH
/protocols/fpolicy/{svm.uuid}/connections/{node.uuid}/{policy.name}/{serve
r}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/9f738ac5-c502-11eb-
b82c-0050568e5902/connections/8ca36b68-c501-11eb-b82c-
0050568e5902/p1/192.168.137.78" -H  "accept: application/json"
```

---

### Retrieve the status of FPolicy servers

GET `/protocols/fpolicy/{svm.uuid}/connections`

**Introduced In:** 9.10

Retrieves the statuses of FPolicy servers.

#### Related ONTAP commands

- `vserver fpolicy show-engine`
- `vserver fpolicy show-passthrough-read-connection`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/connections

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| passthrough_read | boolean | query | False | Whether to view only passthrough-read connections |
| disconnected_reason.message | string | query | False | Filter by disconnected_reason.message |
| disconnected_reason.code | integer | query | False | Filter by disconnected_reason.code |
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| update_time | string | query | False | Filter by update_time |
| svm.name | string | query | False | Filter by svm.name |
| server | string | query | False | Filter by server |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| type | string | query | False | Filter by type |
| policy.name | string | query | False | Filter by policy.name |
| session_uuid | string | query | False | Filter by session_uuid |
| max_records | integer | query | False | Limit the number of records returned. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[fpolicy_connection] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "disconnected_reason": {
        "code": 9307,
        "message": "TCP Connection to FPolicy server failed."
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "server": "10.132.145.20",
      "session_uuid": "5224ec64-b336-11eb-841c-0050568e14c2",
      "state": "connected",
      "svm": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "primary",
      "update_time": "2019-06-12 11:00:16 -0400"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

disconnected_reason

Indicates the reason for FPolicy server disconnection.

| Name | Type | Description |
|---|---|---|
| code | integer | Reason ID for the FPolicy Server disconnection. |
| message | string | FPolicy server reason for disconnection message. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

fpolicy_policy_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The name of the FPolicy Policy name. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fpolicy_connection

Displays the connection status information of the FPolicy server.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| disconnected_reason | disconnected_reason | Indicates the reason for FPolicy server disconnection. |
| node | node | |
| policy | fpolicy_policy_reference | |
| server | string | IP address of the FPolicy server. |
| session_uuid | string | Unique session ID associated with each connection to the FPolicy server and it can be used to identify<br>the established connection. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Specifies the FPolicy server connection state indicating if it is in the connected or disconnected state.<br>The following is a list of the possible states:<br>• connected - Connected<br>• disconnected - Disconnected |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | FPolicy server type. The possible values are:<br>• primary - Primary server<br>• secondary - Secondary server |
| update_time | string | Specifies the time at which FPolicy server is connected or disconnected. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the status of an FPolicy server

GET `/protocols/fpolicy/{svm.uuid}/connections/{node.uuid}/{policy.name}/{server}`

**Introduced In:** 9.10

Retrieves the status of an FPolicy server.

#### Related ONTAP commands

- `vserver fpolicy show-engine`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/connections

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| node.uuid | string | path | True | Node UUID. |
| policy.name | string | path | True | Policy name |
| server | string | path | True | FPolicy Server IP address |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| disconnected_reason | disconnected_reason | Indicates the reason for FPolicy server disconnection. |
| node | node | |
| policy | fpolicy_policy_reference | |
| server | string | IP address of the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| session_uuid | string | Unique session ID associated with each connection to the FPolicy server and it can be used to identify the established connection. |
| state | string | Specifies the FPolicy server connection state indicating if it is in the connected or disconnected state.<br>The following is a list of the possible states:<br>• connected - Connected<br>• disconnected - Disconnected |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | FPolicy server type. The possible values are:<br>• primary - Primary server<br>• secondary - Secondary server |
| update_time | string | Specifies the time at which FPolicy server is connected or disconnected. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "disconnected_reason": {
    "code": 9307,
    "message": "TCP Connection to FPolicy server failed."
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string"
  },
  "server": "10.132.145.20",
  "session_uuid": "5224ec64-b336-11eb-841c-0050568e14c2",
  "state": "connected",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "primary",
  "update_time": "2019-06-12 11:00:16 -0400"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

disconnected_reason

Indicates the reason for FPolicy server disconnection.

| Name | Type | Description |
|---|---|---|
| code | integer | Reason ID for the FPolicy Server disconnection. |
| message | string | FPolicy server reason for disconnection message. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

fpolicy_policy_reference

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The name of the FPolicy Policy name. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the FPolicy server status

PATCH `/protocols/fpolicy/{svm.uuid}/connections/{node.uuid}/{policy.name}/{server}`

**Introduced In:** 9.10

Updates the status of an FPolicy server.

#### Related ONTAP commands

- `vserver fpolicy engine-connect`
- `vserver fpolicy engine-disconnect`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/connections

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| node.uuid | string | path | True | Node UUID. |
| policy.name | string | path | True | Policy name |
| server | string | path | True | FPolicy Server IP address |

#### Request Body

| Name | Type | Description |
|---|---|---|
| disconnected_reason | disconnected_reason | Indicates the reason for FPolicy server disconnection. |
| node | node | |
| server | string | IP address of the FPolicy server. |
| session_uuid | string | Unique session ID associated with each connection to the FPolicy server and it can be used to identify the established connection. |
| state | string | Specifies the FPolicy server connection state indicating if it is in the connected or disconnected state.<br>The following is a list of the possible states:<br>• connected - Connected<br>• disconnected - Disconnected |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | FPolicy server type. The possible values are:<br>• primary - Primary server<br>• secondary - Secondary server |
| update_time | string | Specifies the time at which FPolicy server is connected or disconnected. |

#### Example request

```
{
  "disconnected_reason": {
    "code": 9307,
    "message": "TCP Connection to FPolicy server failed."
  },
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "server": "10.132.145.20",
  "session_uuid": "5224ec64-b336-11eb-841c-0050568e14c2",
  "state": "connected",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "primary",
  "update_time": "2019-06-12 11:00:16 -0400"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 9764954 | The specified policy does not exist |
| 9764911 | Failed to connect to the FPolicy server. Reason: The specified entry does not exist |
| 9764912 | Failed to disconnect the FPolicy server. Reason: The specified entry does not exist |
| 9764948 | The specified policy is disabled. Using a disabled policy with this API is not supported. Use the 'PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}' API to enable the policy |
| 9764963 | The specified server is already connected |
| 9764964 | The specified server is already disconnected |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

disconnected_reason

Indicates the reason for FPolicy server disconnection.

| Name | Type | Description |
|---|---|---|
| code | integer | Reason ID for the FPolicy Server disconnection. |
| message | string | FPolicy server reason for disconnection message. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

self_link

fpolicy_policy_reference

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FPolicy Policy name. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fpolicy_connection

Displays the connection status information of the FPolicy server.

| Name | Type | Description |
| --- | --- | --- |
| disconnected_reason | disconnected_reason | Indicates the reason for FPolicy server disconnection. |
| node | node | |
| server | string | IP address of the FPolicy server. |
| session_uuid | string | Unique session ID associated with each connection to the FPolicy server and it can be used to identify<br>the established connection. |
| state | string | Specifies the FPolicy server connection state indicating if it is in the connected or disconnected state.<br>The following is a list of the possible states:<br>• connected - Connected<br>• disconnected - Disconnected |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | FPolicy server type. The possible values are:<br>• primary - Primary server<br>• secondary - Secondary server |
| update_time | string | Specifies the time at which FPolicy server is connected or disconnected. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FPolicy engine configuration

### Manage FPolicy engine configuration

#### Overview

The FPolicy engine allows you to configure the external servers to which the file access notifications are sent. As part of FPolicy engine configuration, you can configure a server(s) to which the notification is sent, an optional set of secondary server(s) to which the notification is sent in the case of a primary server(s) failure, the port number for FPolicy application, the type of the engine, which is either synchronous or asynchronous and the format of the notifications, which is either xml or protobuf.

For the synchronous engine, ONTAP will wait for a response from the FPolicy application before it allows the operation. With an asynchronous engine, ONTAP proceeds with the operation processing after sending the notification to the FPolicy application. An engine can belong to multiple FPolicy policies. If the format is not specified, the default format, xml, is configured.

#### Examples

### Creating an FPolicy engine

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/engines

#The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"name\": \"engine0\", \"port\": 9876,
\"primary_servers\": [ \"10.132.145.22\", \"10.140.101.109\" ],
\"secondary_servers\": [ \"10.132.145.20\", \"10.132.145.21\" ], \"type\":
\"synchronous\", \"format\": \"xml\", \"request_abort_timeout\": \"PT3M\",
\"request_cancel_timeout\": \"PT29S\",  \"server_progress_timeout\":
\"PT1M\",  \"status_request_interval\": \"PT23S\",
\"keep_alive_interval\": \"PT2M\", \"session_timeout\": \"PT10S\",
\"max_connection_retries\": 5 }"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "engine0",
    "primary_servers": [
      "10.132.145.22",
      "10.140.101.109"
    ],
    "secondary_servers": [
      "10.132.145.20",
      "10.132.145.21"
    ],
    "port": 9876,
    "type": "synchronous",
    "format": "xml"
  }
]
}
```

### Creating an FPolicy engine with the minimum required fields

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/engines

#The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"name\": \"engine0\", \"port\": 9876,
\"primary_servers\": [ \"10.132.145.22\", \"10.140.101.109\" ], \"type\":
\"synchronous\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "engine0",
    "primary_servers": [
      "10.132.145.22",
      "10.140.101.109"
    ],
    "port": 9876,
    "type": "synchronous",
    "format": "xml"
  }
]
}
```

#### Retrieving an FPolicy engine configuration for a particular SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/engines

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
"records": [
  {
```

```
    "svm": {
      "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
    },
    "name": "cifs",
    "primary_servers": [
      "10.20.20.10"
    ],
    "port": 9876,
    "type": "synchronous"
  },
  {
    "svm": {
      "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
    },
    "name": "nfs",
    "primary_servers": [
      "10.23.140.64",
      "10.140.101.109"
    ],
    "secondary_servers": [
      "10.132.145.20",
      "10.132.145.22"
    ],
    "port": 9876,
    "type": "synchronous",
    "format": "xml",
    "ssl_option": "no_auth",
    "resiliency": {
      "enabled": false,
      "retention_duration": "PT3M"
    },
    "buffer_size": {
      "recv_buffer": 262144,
      "send_buffer": 1048576
    },
    "max_server_requests": 500,
    "request_cancel_timeout": "PT29S",
    "request_abort_timeout": "PT3M",
    "server_progress_timeout": "PT1M",
    "status_request_interval": "PT23S",
    "keep_alive_interval": "PT2M",
    "session_timeout": "PT10S",
    "max_connection_retries": 5
  }
],
"num_records": 2
```

```
}
```

### Retrieving a specific FPolicy engine configuration for an SVM

```
# The Api:
GET /protocols/fpolicy/{svm.uuid}/engines/{name}

#The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/cifs?fields=*" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
 },
 "name": "cifs",
 "primary_servers": [
   "10.20.20.10"
 ],
 "port": 9876,
 "type": "synchronous",
 "format": "xml"
}
```

### Updating an FPolicy engine for an SVM

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/engines/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/cifs" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"port\": 6666, \"secondary_servers\": [
\"10.132.145.20\", \"10.132.145.21\" ], \"type\": \"synchronous\"}"
```

### Updating all the attributes of a specific FPolicy engine for an SVM

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/engines/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/cifs" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"port\": 9876, \"primary_servers\": [
\"10.132.145.20\", \"10.140.101.109\" ], \"secondary_servers\": [
\"10.132.145.23\", \"10.132.145.21\" ], \"type\": \"synchronous\",
\"format\": \"protobuf\"}"
```

### Deleting a specific FPolicy engine for an SVM

```
# The API:
DELETE /protocols/fpolicy/{svm.uuid}/engines/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-
11e8-ae49-0050568e2c1e/engines/cifs" -H "accept: application/json"
```

### Retrieve the FPolicy engine configuration for all engines of an SVM

GET `/protocols/fpolicy/{svm.uuid}/engines`

**Introduced In:** 9.6

Retrieves FPolicy engine configurations of all the engines for a specified SVM. ONTAP allows creation of cluster-level FPolicy engines that act as a template for all the SVMs belonging to the cluster. These cluster-level FPolicy engines are also retrieved for the specified SVM.

#### Related ONTAP commands

- `fpolicy policy external-engine show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/engines

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_connection_retries | integer | query | False | Filter by max_connection_retries<br>• Max value: 20<br>• Min value: 0<br>• Introduced in: 9.17 |
| certificate.name | string | query | False | Filter by certificate.name<br>• Introduced in: 9.11 |
| certificate.serial_number | string | query | False | Filter by certificate.serial_number<br>• Introduced in: 9.11 |
| certificate.ca | string | query | False | Filter by certificate.ca<br>• Introduced in: 9.11 |
| session_timeout | string | query | False | Filter by session_timeout<br>• Introduced in: 9.17 |
| secondary_servers | string | query | False | Filter by secondary_servers |
| format | string | query | False | Filter by format<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| request_abort_timeout | string | query | False | Filter by request_abort_timeout<br>• Introduced in: 9.11 |
| request_cancel_timeout | string | query | False | Filter by request_cancel_timeout<br>• Introduced in: 9.11 |
| keep_alive_interval | string | query | False | Filter by keep_alive_interval<br>• Introduced in: 9.13 |
| port | integer | query | False | Filter by port |
| max_server_requests | integer | query | False | Filter by max_server_requests<br>• Max value: 10000<br>• Min value: 1<br>• Introduced in: 9.11 |
| buffer_size.send_buffer | integer | query | False | Filter by buffer_size.send_buffer<br>• Max value: 7895160<br>• Min value: 0<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buffer_size.recv_buffer | integer | query | False | Filter by buffer_size.recv_buffer<br>• Max value: 7895160<br>• Min value: 0<br>• Introduced in: 9.11 |
| status_request_interval | string | query | False | Filter by status_request_interval<br>• Introduced in: 9.11 |
| ssl_option | string | query | False | Filter by ssl_option<br>• Introduced in: 9.11 |
| server_progress_timeout | string | query | False | Filter by server_progress_timeout<br>• Introduced in: 9.11 |
| resiliency.directory_path | string | query | False | Filter by resiliency.directory_path<br>• Introduced in: 9.11 |
| resiliency.retention_duration | string | query | False | Filter by resiliency.retention_duration<br>• Introduced in: 9.11 |
| resiliency.enabled | boolean | query | False | Filter by resiliency.enabled<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| type | string | query | False | Filter by type |
| name | string | query | False | Filter by name |
| primary_servers | string | query | False | Filter by primary_servers |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[fpolicy_engine] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "certificate": {
        "ca": "TASample1",
        "name": "Sample1-FPolicy-Client",
        "serial_number": "8DDE112A114D1FBC"
      },
      "format": "string",
      "keep_alive_interval": "PT2M",
      "max_connection_retries": 5,
      "max_server_requests": 500,
      "name": "fp_ex_eng",
      "port": 9876,
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "request_abort_timeout": "PT40S",
      "request_cancel_timeout": "PT20S",
      "resiliency": {
        "directory_path": "/dir1",
        "retention_duration": "PT3M"
      },
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "server_progress_timeout": "PT1M",
      "session_timeout": "PT10S",
      "ssl_option": "string",
      "status_request_interval": "PT10S",
      "svm": {
        "uuid": "string"
      },
```

```
        "type": "string"
      }
    ]
  }
```

#### Error

```
Status: Default,
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_engine

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |

| Name | Type | Description |
| --- | --- | --- |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create the FPolicy engine configuration for an SVM**

POST `/protocols/fpolicy/{svm.uuid}/engines`

**Introduced In:** 9.6

Creates an FPolicy engine configuration for a specified SVM. FPolicy engine creation is allowed only on data SVMs.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the FPolicy engine.
- `name` - Name of external engine.
- `port` - Port number of the FPolicy server application.
- `primary_servers` - List of primary FPolicy servers to which the node will send notifications.

#### Recommended optional properties

- `secondary_servers` - It is recommended to configure secondary FPolicy server to which the node will send notifications when the primary server is down.

#### Default property values

- `type` - *synchronous*
- `format` - *xml*

#### Related ONTAP commands

- `fpolicy policy external-engine create`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/engines

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |

| Name | Type | Description |
| --- | --- | --- |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br><br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

### Example request

```
{
  "certificate": {
    "ca": "TASample1",
    "name": "Sample1-FPolicy-Client",
    "serial_number": "8DDE112A114D1FBC"
  },
  "format": "string",
  "keep_alive_interval": "PT2M",
  "max_connection_retries": 5,
  "max_server_requests": 500,
  "name": "fp_ex_eng",
  "port": 9876,
  "primary_servers": [
    "10.132.145.20",
    "10.140.101.109"
  ],
  "request_abort_timeout": "PT40S",
  "request_cancel_timeout": "PT20S",
  "resiliency": {
    "directory_path": "/dir1",
    "retention_duration": "PT3M"
  },
  "secondary_servers": [
    "10.132.145.20",
    "10.132.145.21"
  ],
  "server_progress_timeout": "PT1M",
  "session_timeout": "PT10S",
  "ssl_option": "string",
  "status_request_interval": "PT10S",
  "svm": {
    "uuid": "string"
  },
  "type": "string"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[fpolicy_engine] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "certificate": {
        "ca": "TASample1",
        "name": "Sample1-FPolicy-Client",
        "serial_number": "8DDE112A114D1FBC"
      },
      "format": "string",
      "keep_alive_interval": "PT2M",
      "max_connection_retries": 5,
      "max_server_requests": 500,
      "name": "fp_ex_eng",
      "port": 9876,
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "request_abort_timeout": "PT40S",
      "request_cancel_timeout": "PT20S",
      "resiliency": {
        "directory_path": "/dir1",
        "retention_duration": "PT3M"
      },
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "server_progress_timeout": "PT1M",
      "session_timeout": "PT10S",
      "ssl_option": "string",
      "status_request_interval": "PT10S",
      "svm": {
        "uuid": "string"
      },
      "type": "string"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9764885 | The primary secondary server has a redundant IP address |
| 9764953 | The name of the FPolicy engine is "native" which is reserved by the system |
| 9765011 | The resiliency feature is not supported with mandatory screening |
| 9765012 | The specified resiliency directory path does not exist |
| 9765004 | The specified session timeout exceeds the maximum limit |
| 9764935 | The specified max connection retries exceeds the maximum limit |

#### Definitions

**See Definitions**

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_engine

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |

| Name | Type | Description |
| --- | --- | --- |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an FPolicy external engine configuration

DELETE `/protocols/fpolicy/{svm.uuid}/engines/{name}`

**Introduced In:** 9.6

Deletes the FPolicy external engine configuration. Deletion of an FPolicy engine that is attached to one or more FPolicy policies is not allowed.

#### Related ONTAP commands

- `fpolicy policy external-engine delete`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/engines

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9764940 | At least one FPolicy policy is using the FPolicy engine |
| 9764887 | The FPolicy engine is a cluster level FPolicy engine |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a particular FPolicy engine configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}/engines/{name}`

**Introduced In:** 9.6

Retrieves a particular FPolicy engine configuration of a specified SVM. A cluster-level FPolicy engine configuration cannot be retrieved for a data SVM.

#### Related ONTAP commands

- `fpolicy policy external-engine show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/engines

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| name | string | Specifies the name to assign to the external server configuration. |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br><br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

### Example response


```
{
  "certificate": {
    "ca": "TASample1",
    "name": "Sample1-FPolicy-Client",
    "serial_number": "8DDE112A114D1FBC"
  },
  "format": "string",
  "keep_alive_interval": "PT2M",
  "max_connection_retries": 5,
  "max_server_requests": 500,
  "name": "fp_ex_eng",
  "port": 9876,
  "primary_servers": [
    "10.132.145.20",
    "10.140.101.109"
  ],
  "request_abort_timeout": "PT40S",
  "request_cancel_timeout": "PT20S",
  "resiliency": {
    "directory_path": "/dir1",
    "retention_duration": "PT3M"
  },
  "secondary_servers": [
    "10.132.145.20",
    "10.132.145.21"
  ],
  "server_progress_timeout": "PT1M",
  "session_timeout": "PT10S",
  "ssl_option": "string",
  "status_request_interval": "PT10S",
  "svm": {
    "uuid": "string"
  },
  "type": "string"
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a specific FPolicy engine configuration for an SVM

PATCH `/protocols/fpolicy/{svm.uuid}/engines/{name}`

**Introduced In:** 9.6

Updates a specific FPolicy engine configuration of an SVM. Modification of an FPolicy engine that is attached to one or more enabled FPolicy policies is not allowed.

#### Related ONTAP commands

- `fpolicy policy external-engine modify`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/engines

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |

| Name | Type | Description |
| --- | --- | --- |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br><br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

**Example request**

```
{
  "certificate": {
    "ca": "TASample1",
    "name": "Sample1-FPolicy-Client",
    "serial_number": "8DDE112A114D1FBC"
  },
  "format": "string",
  "keep_alive_interval": "PT2M",
  "max_connection_retries": 5,
  "max_server_requests": 500,
  "port": 9876,
  "primary_servers": [
    "10.132.145.20",
    "10.140.101.109"
  ],
  "request_abort_timeout": "PT40S",
  "request_cancel_timeout": "PT20S",
  "resiliency": {
    "directory_path": "/dir1",
    "retention_duration": "PT3M"
  },
  "secondary_servers": [
    "10.132.145.20",
    "10.132.145.21"
  ],
  "server_progress_timeout": "PT1M",
  "session_timeout": "PT10S",
  "ssl_option": "string",
  "status_request_interval": "PT10S",
  "svm": {
    "uuid": "string"
  },
  "type": "string"
}
```

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9764922 | The primary and secondary server has a redundant IP address |
| 9764942 | At least one FPolicy policy is using the FPolicy engine |
| 9764886 | FPolicy engine is a cluster-level FPolicy engine |
| 9765011 | The resiliency feature is not supported with mandatory screening |
| 9765012 | The specified resiliency directory path does not exist |
| 9765042 | The specified send buffer size exceeds the maximum limit |
| 9765043 | The specified receive buffer size exceeds the maximum limit |
| 9765063 | Policy with Persistent Store feature does not support a "synchronous" |
| 9765004 | The specified session timeout exceeds the maximum limit |
| 9764935 | The specified max connection retries exceeds the maximum limit |

#### Definitions

**See Definitions**

buffer_size

Specifies the send and receive buffer size of the connected socket for the FPolicy server.

| Name | Type | Description |
|---|---|---|
| recv_buffer | integer | Specifies the receive buffer size of the connected socket for the FPolicy server. Default value is 256KB. |
| send_buffer | integer | Specifies the send buffer size of the connected socket for the FPolicy server. Default value 1MB. |

certificate

Provides details about certificate used to authenticate the FPolicy server.

| Name | Type | Description |
|---|---|---|
| ca | string | Specifies the certificate authority (CA) name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |
| name | string | Specifies the certificate name as a fully qualified domain name (FQDN) or custom common name. The certificate is used if SSL authentication between the SVM and the FPolicy server is configured. |
| serial_number | string | Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured. |

resiliency

If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path.

| Name | Type | Description |
| --- | --- | --- |
| directory_path | string | Specifies the directory path under the SVM namespace, where notifications are stored in the files whenever a network outage happens. |
| enabled | boolean | Specifies whether the resiliency feature is enabled or not. Default is false. |
| retention_duration | string | Specifies the ISO-8601 duration, for which the notifications are written to files inside the storage controller during a network outage. The value for this field must be between 0 and 600 seconds. Default is 180 seconds. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_engine

Defines how ONTAP makes and manages connections to external FPolicy servers.

| Name | Type | Description |
| --- | --- | --- |
| buffer_size | buffer_size | Specifies the send and receive buffer size of the connected socket for the FPolicy server. |
| certificate | certificate | Provides details about certificate used to authenticate the FPolicy server. |

| Name | Type | Description |
| --- | --- | --- |
| format | string | The format for the notification messages sent to the FPolicy servers.<br>The possible values are:<br>• xml - Notifications sent to the FPolicy server will be formatted using the XML schema.<br>• protobuf - Notifications sent to the FPolicy server will be formatted using Protobuf schema, which is a binary form. |
| keep_alive_interval | string | Specifies the ISO-8601 interval time for a storage appliance to send Keep Alive message to an FPolicy server. The allowed range is between 10 to 600 seconds. |
| max_connection_retries | integer | This parameter specifies the maximum number of attempts to reconnect to the FPolicy server from an SVM. It is used to specify the number of times a broken connection will be retried. The value for this field must be between 0 and 20. By default, it is 5. |
| max_server_requests | integer | Specifies the maximum number of outstanding requests for the FPolicy server. It is used to specify maximum outstanding requests that will be queued up for the FPolicy server. The value for this field must be between 1 and 10000. The default values are 500, 1000 or 2000 for Low-end(<64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB memory) Platforms respectively. |
| port | integer | Port number of the FPolicy server application. |
| primary_servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| request_abort_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be aborted by a storage appliance. The allowed range is between 0 to 200 seconds. |
| request_cancel_timeout | string | Specifies the ISO-8601 timeout duration for a screen request to be processed by an FPolicy server. The allowed range is between 0 to 100 seconds. |
| resiliency | resiliency | If all primary and secondary servers are down, or if no response is received from the FPolicy servers, file access events are stored inside the storage controller under the specified resiliency-directory-path. |
| secondary_servers | array[string] | |
| server_progress_timeout | string | Specifies the ISO-8601 timeout duration in which a throttled FPolicy server must complete at least one screen request. If no request is processed within the timeout, connection to the FPolicy server is terminated. The allowed range is between 0 to 100 seconds. |
| session_timeout | string | This parameter specifies the interval after which a new session ID is sent to the FPolicy server during reconnection attempts. The default value is set to 10 seconds. If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the -session-timeout interval, the old session ID is sent to the FPolicy server so that it can send responses for old notifications. |

| Name | Type | Description |
| --- | --- | --- |
| ssl_option | string | Specifies the SSL option for external communication with the FPolicy server. Possible values include the following:<br>• no_auth When set to "no_auth", no authentication takes place.<br>• server_auth When set to "server_auth", only the FPolicy server is authenticated by the SVM. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.<br>• mutual_auth When set to "mutual_auth", mutual authentication takes place between the SVM and the FPolicy server. This means authentication of the FPolicy server by the SVM along with authentication of the SVM by the FPolicy server. With this option, before creating the FPolicy external engine, the administrator must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. |
| status_request_interval | string | Specifies the ISO-8601 interval time for a storage appliance to query a status request from an FPolicy server. The allowed range is between 0 to 50 seconds. |
| svm | svm | |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:<br>• synchronous - After sending a notification, wait for a response from the FPolicy server.<br>• asynchronous - After sending a notification, file request processing continues.<br>◦ Default value: 1<br>◦ enum: ["synchronous", "asynchronous"]<br>◦ Introduced in: 9.6<br>◦ x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FPolicy event configuration

### Manage FPolicy event configuration

#### Overview

FPolicy events configurations allow you to specify which file access is monitored. As part of an FPolicy event, you can configure the SVM for which the events are generated, the name of the event configuration, the protocol (cifs, nfsv3/nfsv4) for which the events are generated, the file operations which are monitored, and filters that can be used to filter the unwanted notification generation for a specified protocol and file operation.

Each protocol has a set of supported file operations and filters. An SVM can have multiple events. A single FPolicy policy can have multiple FPolicy events.

FPolicy events can also be configured to monitor file operations which fail due to lack of permissions. You can specify which file operation to monitor for failure. However, filters can not be used to filter failed file operations.

#### Examples

### Creating an FPolicy event for a CIFS protocol with all the supported file operations and filters

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/events?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"file_operations\": { \"close\": true, \"create\": true, \"create_dir\":
true, \"delete\": true, \"delete_dir\": true, \"getattr\": true, \"open\":
true, \"read\": true, \"rename\": true, \"rename_dir\": true, \"setattr\":
true, \"write\": true }, \"filters\": { \"close_with_modification\": true,
\"close_with_read\": true, \"close_without_modification\": true,
\"first_read\": true, \"first_write\": true, \"monitor_ads\": true,
\"offline_bit\": true, \"open_with_delete_intent\": true,
\"open_with_write_intent\": true, \"write_with_size_change\": true },
\"name\": \"event_cifs\", \"protocol\": \"cifs\", \"volume_monitoring\":
true}"

# The response:
{
"num_records": 1,
  "records": [
    {
      "name": "event_cifs",
      "protocol": "cifs",
      "volume_monitoring": true,
      "monitor_fileop_failure": false,
      "file_operations": {
        "close": true,
        "create": true,
```

```
        "create_dir": true,
        "delete": true,
        "delete_dir": true,
        "getattr": true,
        "open": true,
        "read": true,
        "write": true,
        "rename": true,
        "rename_dir": true,
        "setattr": true
      },
      "filters": {
        "monitor_ads": true,
        "close_with_modification": true,
        "close_without_modification": true,
        "close_with_read": true,
        "first_read": true,
        "first_write": true,
        "offline_bit": true,
        "open_with_delete_intent": true,
        "open_with_write_intent": true,
        "write_with_size_change": true
      }
    }
  ]
}
```


### Creating an FPolicy event for an NFS protocol with all the supported file operations and filters

```
# The API:
post /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/events?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"file_operations\": { \"create\": true, \"create_dir\": true, \"delete\":
true, \"delete_dir\": true, \"link\": true, \"lookup\": true, \"read\":
true, \"rename\": true, \"rename_dir\": true, \"setattr\": true,
\"symlink\": true, \"write\": true }, \"filters\": { \"offline_bit\":
true, \"write_with_size_change\": true }, \"name\": \"event_nfsv3\",
\"protocol\": \"nfsv3\", \"volume_monitoring\": false}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "event_nfsv3",
    "protocol": "nfsv3",
    "volume_monitoring": false,
    "monitor_fileop_failure": false,
    "file_operations": {
      "create": true,
      "create_dir": true,
      "delete": true,
      "delete_dir": true,
      "link": true,
      "lookup": true,
      "read": true,
      "write": true,
      "rename": true,
      "rename_dir": true,
      "setattr": true,
      "symlink": true
    },
    "filters": {
    "offline_bit": true,
    "write_with_size_change": true
    }
  }
]
}
```

**Creating an FPolicy event to monitor failed file operations for an NFS protocol with all the supported file operations**

```
# The API:
post /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/b5087518-40b3-11ed-
b3eb-005056bbe901/events?return_records=false" -H "accept:
application/json" -H  "authorization: Basic <CREDENTIALS>" -H  "Content-
Type: application/json" -d "{  \"file_operations\": {    \"create\": true,
\"create_dir\": true,    \"delete\": true,    \"delete_dir\": true,
\"link\": true,    \"read\": true,    \"rename\": true,    \"rename_dir\":
true,    \"write\": true  },  \"name\": \"nfs_failed_op\",  \"protocol\":
\"nfsv3\",  \"monitor_fileop_failure\": true,  \"volume_monitoring\":
false}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "nfs_failed_op",
    "protocol": "nfsv3",
    "volume_monitoring": false,
    "monitor_fileop_failure": true,
    "file_operations": {
      "create": true,
      "create_dir": true,
      "delete": true,
      "delete_dir": true,
      "link": true,
      "read": true,
      "write": true,
      "rename": true,
      "rename_dir": true
    }
  }
]
}
```

**Retrieving all of the FPolicy event configurations configured to monitor failed file operations for a specified SVM**

```
# The API:
post /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/b5087518-40b3-11ed-
b3eb-
005056bbe901/events?monitor_fileop_failure=true&fields=*&return_records=tr
ue&return_timeout=15" -H  "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
    "uuid": "b5087518-40b3-11ed-b3eb-005056bbe901"
    },
    "name": "fo_event",
    "protocol": "cifs",
    "volume_monitoring": false,
    "monitor_fileop_failure": true,
    "file_operations": {
      "close": false,
      "create": false,
      "create_dir": false,
      "delete": false,
      "delete_dir": false,
      "getattr": false,
      "link": false,
      "lookup": false,
      "open": true,
      "read": false,
      "write": false,
      "rename": false,
      "rename_dir": false,
      "setattr": false,
      "symlink": false
    }
  }
],
"num_records": 1
}
```

**Retrieving all of the FPolicy event configurations for a specified SVM**

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/events/?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
    },
    "name": "cluster",
    "protocol": "cifs",
    "volume_monitoring": false,
    "monitor_fileop_failure": false,
    "file_operations": {
      "close": true,
      "create": false,
      "create_dir": false,
      "delete": false,
      "delete_dir": false,
      "getattr": false,
      "link": false,
      "lookup": false,
      "open": false,
      "read": false,
      "write": false,
      "rename": false,
      "rename_dir": false,
      "setattr": false,
      "symlink": false
    },
    "filters": {
      "monitor_ads": false,
      "close_with_modification": false,
      "close_without_modification": false,
      "close_with_read": true,
      "first_read": false,
```

```
      "first_write": false,
      "offline_bit": false,
      "open_with_delete_intent": false,
      "open_with_write_intent": false,
      "write_with_size_change": false,
      "setattr_with_owner_change": false,
      "setattr_with_group_change": false,
      "setattr_with_sacl_change": false,
      "setattr_with_dacl_change": false,
      "setattr_with_modify_time_change": false,
      "setattr_with_access_time_change": false,
      "setattr_with_creation_time_change": false,
      "setattr_with_mode_change": false,
      "setattr_with_size_change": false,
      "setattr_with_allocation_size_change": false,
      "exclude_directory": false
    }
  },
  {
    "svm": {
      "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
    },
    "name": "event_cifs",
    "protocol": "cifs",
    "volume_monitoring": true,
    "monitor_fileop_failure": false,
    "file_operations": {
      "close": true,
      "create": true,
      "create_dir": true,
      "delete": true,
      "delete_dir": true,
      "getattr": true,
      "link": false,
      "lookup": false,
      "open": true,
      "read": true,
      "write": true,
      "rename": true,
      "rename_dir": true,
      "setattr": true,
      "symlink": false
    },
    "filters": {
      "monitor_ads": true,
      "close_with_modification": true,
```

```
      "close_without_modification": true,
      "close_with_read": true,
      "first_read": true,
      "first_write": true,
      "offline_bit": true,
      "open_with_delete_intent": true,
      "open_with_write_intent": true,
      "write_with_size_change": true,
      "setattr_with_owner_change": false,
      "setattr_with_group_change": false,
      "setattr_with_sacl_change": false,
      "setattr_with_dacl_change": false,
      "setattr_with_modify_time_change": false,
      "setattr_with_access_time_change": false,
      "setattr_with_creation_time_change": false,
      "setattr_with_mode_change": false,
      "setattr_with_size_change": false,
      "setattr_with_allocation_size_change": false,
      "exclude_directory": false
    }
  }
],
"num_records": 2
}
```


#### Retrieving a specific FPolicy event configuration for an SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/events/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-
0050568e2c1e/events/event_cifs?fields=*&return_records=true&return_timeout
=15" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
},
"name": "event_cifs",
```

```
"protocol": "cifs",
"volume_monitoring": true,
"monitor_fileop_failure": false,
"file_operations": {
  "close": true,
  "create": true,
  "create_dir": true,
  "delete": true,
  "delete_dir": true,
  "getattr": true,
  "link": false,
  "lookup": false,
  "open": true,
  "read": true,
  "write": true,
  "rename": true,
  "rename_dir": true,
  "setattr": true,
  "symlink": false
},
"filters": {
  "monitor_ads": true,
  "close_with_modification": true,
  "close_without_modification": true,
  "close_with_read": true,
  "first_read": true,
  "first_write": true,
  "offline_bit": true,
  "open_with_delete_intent": true,
  "open_with_write_intent": true,
  "write_with_size_change": true,
  "setattr_with_owner_change": false,
  "setattr_with_group_change": false,
  "setattr_with_sacl_change": false,
  "setattr_with_dacl_change": false,
  "setattr_with_modify_time_change": false,
  "setattr_with_access_time_change": false,
  "setattr_with_creation_time_change": false,
  "setattr_with_mode_change": false,
  "setattr_with_size_change": false,
  "setattr_with_allocation_size_change": false,
  "exclude_directory": false
}
}
```

### Updating a specific FPolicy event configuration for a specified SVM

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/events/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/events/event_cifs" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"file_operations\": { \"close\":
false, \"create\": false, \"read\": true }, \"filters\": {
\"close_with_modification\": false, \"close_with_read\": false,
\"close_without_modification\": false }, \"protocol\": \"cifs\",
\"volume_monitoring\": false}"
```

### Deleting a specific FPolicy event configuration for a specific SVM

```
# The API:
DELETE /protocols/fpolicy/{svm.uuid}/events/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-
11e8-ae49-0050568e2c1e/events/event_cifs" -H "accept: application/json"
```

### Retrieve the FPolicy event configuration for all events for an SVM

GET `/protocols/fpolicy/{svm.uuid}/events`

**Introduced In:** 9.6

Retrieves FPolicy event configurations for all events for a specified SVM. ONTAP allows the creation of cluster-level FPolicy events that act as a template for all the data SVMs belonging to the cluster. These cluster-level FPolicy events are also retrieved for the specified SVM.

#### Related ONTAP commands

- `fpolicy policy event show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/events

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol | string | query | False | Filter by protocol |
| filters.first_write | boolean | query | False | Filter by filters.first_write |
| filters.setattr_with_modify_time_change | boolean | query | False | Filter by filters.setattr_with_modify_time_change |
| filters.monitor_ads | boolean | query | False | Filter by filters.monitor_ads |
| filters.setattr_with_dacl_change | boolean | query | False | Filter by filters.setattr_with_dacl_change |
| filters.close_with_read | boolean | query | False | Filter by filters.close_with_read |
| filters.write_with_size_change | boolean | query | False | Filter by filters.write_with_size_change |
| filters.setattr_with_owner_change | boolean | query | False | Filter by filters.setattr_with_owner_change |
| filters.exclude_directory | boolean | query | False | Filter by filters.exclude_directory |
| filters.first_read | boolean | query | False | Filter by filters.first_read |
| filters.open_with_delete_intent | boolean | query | False | Filter by filters.open_with_delete_intent |
| filters.close_without_modification | boolean | query | False | Filter by filters.close_without_modification |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| filters.setattr_with_size_change | boolean | query | False | Filter by filters.setattr_with_size_change |
| filters.setattr_with_access_time_change | boolean | query | False | Filter by filters.setattr_with_access_time_change |
| filters.setattr_with_sacl_change | boolean | query | False | Filter by filters.setattr_with_sacl_change |
| filters.setattr_with_mode_change | boolean | query | False | Filter by filters.setattr_with_mode_change |
| filters.setattr_with_group_change | boolean | query | False | Filter by filters.setattr_with_group_change |
| filters.setattr_with_allocation_size_change | boolean | query | False | Filter by filters.setattr_with_allocation_size_change |
| filters.open_with_write_intent | boolean | query | False | Filter by filters.open_with_write_intent |
| filters.close_with_modification | boolean | query | False | Filter by filters.close_with_modification |
| filters.setattr_with_creation_time_change | boolean | query | False | Filter by filters.setattr_with_creation_time_change |
| filters.offline_bit | boolean | query | False | Filter by filters.offline_bit |
| name | string | query | False | Filter by name |
| volume_monitoring | boolean | query | False | Filter by volume_monitoring |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| monitor_fileop_failure | boolean | query | False | Filter by monitor_fileop_failure<br>• Introduced in: 9.13 |
| file_operations.read | boolean | query | False | Filter by file_operations.read |
| file_operations.write | boolean | query | False | Filter by file_operations.write |
| file_operations.getattr | boolean | query | False | Filter by file_operations.getattr |
| file_operations.create | boolean | query | False | Filter by file_operations.create |
| file_operations.close | boolean | query | False | Filter by file_operations.close |
| file_operations.setattr | boolean | query | False | Filter by file_operations.setattr |
| file_operations.lookup | boolean | query | False | Filter by file_operations.lookup |
| file_operations.create_dir | boolean | query | False | Filter by file_operations.create_dir |
| file_operations.delete | boolean | query | False | Filter by file_operations.delete |
| file_operations.access | boolean | query | False | Filter by file_operations.access<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| file_operations.open | boolean | query | False | Filter by file_operations.open |
| file_operations.delete_dir | boolean | query | False | Filter by file_operations.delete_dir |
| file_operations.rename_dir | boolean | query | False | Filter by file_operations.rename_dir |
| file_operations.symlink | boolean | query | False | Filter by file_operations.symlink |
| file_operations.rename | boolean | query | False | Filter by file_operations.rename |
| file_operations.link | boolean | query | False | Filter by file_operations.link |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[fpolicy_event] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "name": "event_cifs",
      "protocol": "string",
      "svm": {
        "uuid": "string"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
|---|---|---|
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |

| Name | Type | Description |
| --- | --- | --- |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |

| Name | Type | Description |
| --- | --- | --- |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the FPolicy event configuration for an SVM

POST `/protocols/fpolicy/{svm.uuid}/events`

**Introduced In:** 9.6

Creates an FPolicy event configuration for a specified SVM. FPolicy event creation is allowed only on data SVMs. When a protocol is specified, you must specify a file operation or a file operation and filters. When FPolicy event is configured to monitor failed file operations, you must specify protocol and file operations. Filters are not supported when failed file operations are monitored.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the FPolicy event.
- `name` - Name of the FPolicy event.

#### Recommended optional properties

- `file-operations` - List of file operations to monitor.
- `protocol` - Protocol for which the file operations should be monitored.
- `filters` - List of filters for the specified file operations.
- `monitor-fileop-failure` - Enabled monitoring of failed file operations.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `file_operations.*` - *false*

- `filters.*` - *false*
- `volume-monitoring` - *false*
- `monitor-fileop-failure` - *false*

#### Related ONTAP commands

- `fpolicy policy event create`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/events

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters. The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

**Example request**

```
{
  "name": "event_cifs",
  "protocol": "string",
  "svm": {
    "uuid": "string"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[fpolicy_event] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "name": "event_cifs",
      "protocol": "string",
      "svm": {
        "uuid": "string"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

| Error Code | Description |
|---|---|
| 9764929 | The file operation is not supported by the protocol |
| 9764955 | The filter is not supported by the protocol |
| 9764930 | The filter is not supported by any of the file operations |
| 9764946 | The protocol is specified without a file operation or a file operation and filter pair |

#### Definitions

**See Definitions**

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
|---|---|---|
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |

| Name | Type | Description |
| --- | --- | --- |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
| --- | --- | --- |
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a specific FPolicy event configuration for an SVM

DELETE `/protocols/fpolicy/{svm.uuid}/events/{name}`

**Introduced In:** 9.6

Deletes a specific FPolicy event configuration for an SVM. A cluster-level FPolicy event configuration cannot be modified for a data SVM through REST. An FPolicy event that is attached to an FPolicy policy cannot be deleted.

#### Related ONTAP commands

- `fpolicy policy event delete`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/events

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

| Error Code | Description |
|---|---|
| 9764874 | The FPolicy event is a cluster event |
| 9764947 | The FPolicy event is attached to an FPolicy policy |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a specific FPolicy event configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}/events/{name}`

**Introduced In:** 9.6

Retrieves a specific FPolicy event configuration for an SVM. A cluster-level FPolicy event configuration cannot be retrieved for a data SVM through a REST API.

#### Related ONTAP commands

- `fpolicy policy event show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/events

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| name | string | Specifies the name of the FPolicy event. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

### Example response

```
{
  "name": "event_cifs",
  "protocol": "string",
  "svm": {
    "uuid": "string"
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
|---|---|---|
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |

| Name | Type | Description |
| --- | --- | --- |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a specific FPolicy event configuration for an SVM

PATCH `/protocols/fpolicy/{svm.uuid}/events/{name}`

**Introduced In:** 9.6

Updates a specific FPolicy event configuration for an SVM. A cluster-level FPolicy event configuration cannot be modified for a data SVM through REST. When the file operations and filters fields are modified, the previous values are retained and new values are added to the list of previous values. To remove a particular file operation or filter, set its value to false in the request.

#### Related ONTAP commands

- `fpolicy policy event modify`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/events

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

**Example request**

```
{
  "protocol": "string",
  "svm": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

| Error Code | Description |
|---|---|
| 9764873 | The event is a cluster event |
| 9764929 | The file operation is not supported by the protocol |
| 9764955 | The filter is not supported by the protocol |
| 9764930 | The filter is not supported by any of the file operations |
| 9764946 | The protocol is specified without file operation or a file operation and filter pair |
| 9765048 | The monitor fileop failure option is set without protocol and file operations |

#### Definitions

**See Definitions**

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

| Name | Type | Description |
|---|---|---|
| access | boolean | Access operations |
| close | boolean | File close operations |
| create | boolean | File create operations |
| create_dir | boolean | Directory create operations |
| delete | boolean | File delete operations |
| delete_dir | boolean | Directory delete operations |
| getattr | boolean | Get attribute operations |
| link | boolean | Link operations |
| lookup | boolean | Lookup operations |
| open | boolean | File open operations |
| read | boolean | File read operations |
| rename | boolean | File rename operations |
| rename_dir | boolean | Directory rename operations |
| setattr | boolean | Set attribute operations |
| symlink | boolean | Symbolic link operations |
| write | boolean | File write operations |

filters

Specifies the list of filters for a given file operation for the specified protocol.
When you specify the filters, you must specify the valid protocols and a valid file operations.

| Name | Type | Description |
| --- | --- | --- |
| close_with_modification | boolean | Filter the client request for close with modification. |
| close_with_read | boolean | Filter the client request for close with read. |
| close_without_modification | boolean | Filter the client request for close without modification. |
| exclude_directory | boolean | Filter the client requests for directory operations. When this filter is specified directory operations are not monitored. |
| first_read | boolean | Filter the client requests for the first-read. |
| first_write | boolean | Filter the client requests for the first-write. |
| monitor_ads | boolean | Filter the client request for alternate data stream. |
| offline_bit | boolean | Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed. |
| open_with_delete_intent | boolean | Filter the client request for open with delete intent. |
| open_with_write_intent | boolean | Filter the client request for open with write intent. |
| setattr_with_access_time_change | boolean | Filter the client setattr requests for changing the access time of a file or directory. |
| setattr_with_allocation_size_change | boolean | Filter the client setattr requests for changing the allocation size of a file. |
| setattr_with_creation_time_change | boolean | Filter the client setattr requests for changing the creation time of a file or directory. |

| Name | Type | Description |
| --- | --- | --- |
| setattr_with_dacl_change | boolean | Filter the client setattr requests for changing dacl on a file or directory. |
| setattr_with_group_change | boolean | Filter the client setattr requests for changing group of a file or directory. |
| setattr_with_mode_change | boolean | Filter the client setattr requests for changing the mode bits on a file or directory. |
| setattr_with_modify_time_change | boolean | Filter the client setattr requests for changing the modification time of a file or directory. |
| setattr_with_owner_change | boolean | Filter the client setattr requests for changing owner of a file or directory. |
| setattr_with_sacl_change | boolean | Filter the client setattr requests for changing sacl on a file or directory. |
| setattr_with_size_change | boolean | Filter the client setattr requests for changing the size of a file. |
| write_with_size_change | boolean | Filter the client request for write with size change. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

| Name | Type | Description |
| --- | --- | --- |
| file_operations | file_operations | Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.<br>The event will check the operations specified from all client requests using the protocol. |
| filters | filters | Specifies the list of filters for a given file operation for the specified protocol.<br>When you specify the filters, you must specify the valid protocols and a valid file operations. |
| monitor_fileop_failure | boolean | Specifies whether failed file operations monitoring is required. |
| protocol | string | Protocol for which event is created. If you specify protocol, then you<br>must also specify a valid value for the file operation parameters.<br>The value of this parameter must be one of the following:<br>• cifs - for the CIFS protocol.<br>• nfsv3 - for the NFSv3 protocol.<br>• nfsv4 - for the NFSv4 protocol. |
| svm | svm | |
| volume_monitoring | boolean | Specifies whether volume operation monitoring is required. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FPolicy persistent store configuration

### Manage FPolicy persistent store configuration

#### Overview

- The FPolicy Persistent Store feature enables the configuration of a Persistent Store. This includes:
- setting up a store to retain event notifications
- specifying the volume created for FPolicy Persistent Store
- defining the size of the volume
- defining the autosize mode for the volume.
- Each SVM can only have one Persistent Store. The same Persistent Store can be used by multiple policies within the same SVM. Once a Persistent Store is created, it can be utilized in the FPolicy policy configuration for the async and non-mandatory engine.

#### Examples

### Creating an FPolicy Persistent Store with all required parameters

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/persistent-stores

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/persistent-stores?return_records=true" -H "accept:
application/json" -d "{\"name\": \"ps1\", \"volume\": \"psvol\", \"size\":
1073741824, \"autosize_mode\": \"off\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "ps1",
    "volume": "psvol",
    "size": 1073741824,
    "autosize_mode": "off"
  }
]
}
```

### Creating an FPolicy Persistent Store with the minimum required fields

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/persistent-stores

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/persistent-stores?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{\"name\":
\"ps1\", \"volume\": \"psvol\", \"size\": 1073741824 }"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "ps1",
    "volume": "psvol",
    "size": 1073741824
  }
]
}
```

#### Retrieving an FPolicy Persistent Store configuration for a specific SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/persistent-stores

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/persistent-
stores/?return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "name": "ps1"
  }
],
"num_records": 1
}
```

### Retrieving a specific FPolicy Persistent Store configuration for a specific SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/persistent-stores/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/persistent-
stores/persistent_fpolicy?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
},
"name": "ps1",
"volume": "psvol",
"size": 1073741824,
"autosize_mode": "off"
}
```

### Updating an FPolicy Persistent Store for an SVM

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/persistent-stores/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/persistent-stores/persistent_fpolicy" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"volume\":
\"psvol\" , \"size\": 1073741824}"
```

### Deleting a specific FPolicy Persistent Store configuration for a specific SVM

```
# The API:
DELETE /protocols/fpolicy/{svm.uuid}/persistent-stores/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-
11e8-ae49-0050568e2c1e/persistent-stores/persistent_fpolicy" -H "accept:
application/json"
```

---

### Retrieve the FPolicy persistent store configurations for an SVM

GET `/protocols/fpolicy/{svm.uuid}/persistent-stores`

**Introduced In:** 9.14

Retrieves FPolicy Persistent Store configurations for a specified SVM.

#### Related ONTAP commands

- `fpolicy persistent-store show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/persistent-stores

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume | string | query | False | Filter by volume |
| name | string | query | False | Filter by name |
| size | integer | query | False | Filter by size<br>• Introduced in: 9.15 |
| autosize_mode | string | query | False | Filter by autosize_mode<br>• Introduced in: 9.15 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[fpolicy_persistent_store] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "autosize_mode": "string",
      "name": "ps1",
      "size": "100M",
      "svm": {
        "uuid": "string"
      },
      "volume": "psvol"
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

svm

| Name | Type | Description |
|---|---|---|
| uuid | string | SVM UUID |

fpolicy_persistent_store

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
|---|---|---|
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an FPolicy persistent store configuration for an SVM

POST `/protocols/fpolicy/{svm.uuid}/persistent-stores`

**Introduced In:** 9.14

Creates an FPolicy Persistent Store configuration for a specified SVM.

Important notes:

- FPolicy Persistent Store creation is allowed only on data SVMs.
- In persistent mode, when the Persistent Store is full, event notifications are dropped.
- There is flexibility to provide an existing volume or create a new volume for the persistent storage. The creation of new volume is handled internally.
- For existing volumes, the snapshot policy is set to 'none' and the size is adjusted to the specified value.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the FPolicy Persistent Store.
- `name` - Name of the FPolicy Persistent Store.
- `volume` - Volume specified for Persistent Store (only FlexVol volumes of type RW are supported).

#### Optional properties

- `size` - Size of the Persistent Store volume.
- `autosize-mode` - Autosize mode for the Persistent Store volume.

#### Related ONTAP commands

- `fpolicy persistent-store create`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/persistent-stores

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

### Example request

```
{
  "autosize_mode": "string",
  "name": "ps1",
  "size": "100M",
  "svm": {
    "uuid": "string"
  },
  "volume": "psvol"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[fpolicy_persistent_store] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "autosize_mode": "string",
      "name": "ps1",
      "size": "100M",
      "svm": {
        "uuid": "string"
      },
      "volume": "psvol"
    }
  ]
}
```

### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

| Error Code | Description |
| --- | --- |
| 9765052 | Cannot create more than one Persistent Store per SVM |
| 9765053 | Volume mentioned does not exist in the SVM |
| 9765068 | The Persistent Store name is too long. It must not exceed 500 characters |
| 2621516 | This operation is only supported on a data SVM |
| 9765069 | Volume is not a FlexVol volume. Only FlexVol volumes are supported for this operation. |
| 9765072 | Volume is not of type RW. Only volumes of type RW are supported. |
| 9765074 | Size is a required parameter for the creation of the Persistent Store. |
| 9765077 | The SVM is not associated with any aggregates. |
| 524849 | Attribute cannot be modified for the target volume because of the specified reason. |
| 917533 | The target volume is not eligible for modification due to certain constraints, such as the volume being offline or currently engaged in a volume move operation. |
| 917826 | Volume minimum autosize must be greater than zero. |
| 918647 | The designated size exceeds the maximum available capacity of the intended volume. |
| 917534 | The value specified for the field \"-size\" is too small. Update the field \"-size\" with the minimum size allowed and retry the operation. |

#### Definitions

**See Definitions**

svm

| Name | Type | Description |
|---|---|---|
| uuid | string | SVM UUID |

fpolicy_persistent_store

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
|---|---|---|
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an FPolicy persistent store configuration for an SVM

DELETE `/protocols/fpolicy/{svm.uuid}/persistent-stores/{name}`

**Introduced In:** 9.14

Deletes a specific FPolicy Persistent Store configuration for an SVM.

Important notes:
An option has been provided to delete the associated volume.

#### Related ONTAP commands

- `fpolicy persistent-store delete`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/persistent-stores

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| delete_volume | boolean | query | False | Delete the associated volume of the Persistent Store.<br>• Introduced in: 9.15<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

| Error Code | Description |
|---|---|
| 9765064 | Persistent Store is active on at least one policy |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FPolicy persistent store configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}/persistent-stores/{name}`

**Introduced In:** 9.14

Retrieves a specific FPolicy Persistent Store configuration for an SVM.

#### Related ONTAP commands

- `fpolicy persistent-store show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/persistent-stores

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| autosize_mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows in response to the amount of space used.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |
| name | string | The name specified for the FPolicy Persistent Store. |
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

#### Example response

```
{
  "autosize_mode": "string",
  "name": "ps1",
  "size": "100M",
  "svm": {
    "uuid": "string"
  },
  "volume": "psvol"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

svm

| Name | Type | Description |
|---|---|---|
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an FPolicy persistent store configuration for an SVM

PATCH `/protocols/fpolicy/{svm.uuid}/persistent-stores/{name}`

**Introduced In:** 9.14

Updates a specific FPolicy Persistent Store configuration for an SVM.

Important notes:

- If the volume exists, it is set to the size specified and the snapshot policy is set to "none". Otherwise, a new volume is created.
- The autosize parameter is not available in PATCH operations for this endpoint, use the the "autosize" parameter in PATCH for the "/storage/volumes/{uuid}" endpoint instead.
- When the Persistent Store is updated with a new volume, the previous volume is not automatically deleted. An option is provided to delete the previous volume.

#### Related ONTAP commands

- `fpolicy persistent-store modify`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/persistent-stores

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| delete_volume | boolean | query | False | Delete the previous volume of a Persistent Store if a new volume is provided.<br>• Introduced in: 9.15<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

### Example request

```
{
  "size": "100M",
  "svm": {
    "uuid": "string"
  },
  "volume": "psvol"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

| Error Code | Description |
|---|---|
| 9765053 | Volume mentioned does not exist in the SVM |
| 9765056 | Specified Persistent Store name does not exist in the SVM |
| 9765069 | Volume is not a FlexVol volume. Only FlexVol volumes are supported for this operation. |
| 9765070 | Cannot update a Persistent Store volume when it is part of an enabled policy |
| 9765072 | Volume is not of type RW. Only volumes of type RW are supported. |
| 9765074 | Size is a required parameter for the creation of the Persistent Store. |
| 9765077 | The SVM is not associated with any aggregates. |
| 524849 | Attribute cannot be modified for the target volume because of the specified reason. |
| 917533 | The target volume is not eligible for modification due to certain constraints, such as the volume being offline or currently engaged in a volume move operation. |
| 917826 | Volume minimum autosize must be greater than zero. |
| 918647 | The designated size exceeds the maximum available capacity of the intended volume. |

| Error Code | Description |
| --- | --- |
| 917534 | The value specified for the field \"-size\" is too small. Update the field \"-size\" with the minimum size allowed and retry the operation. |

**Definitions**

**See Definitions**

svm

| Name | Type | Description |
|---|---|---|
| uuid | string | SVM UUID |

fpolicy_persistent_store

The information that an FPolicy process needs in order to configure a Persistent Store.

| Name | Type | Description |
|---|---|---|
| size | integer | The size of the Persistent Store volume, in bytes. |
| svm | svm | |
| volume | string | The specified volume to store the events for the FPolicy Persistent Store. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SVM FPolicy configuration

## Manage SVM FPolicy configuration

### Overview

The FPolicy policy acts as a container for different constituents of the FPolicy such as FPolicy events and the FPolicy engine. It also provides a platform for policy management functions, such as policy enabling and disabling. As part of FPolicy policy configuration, you can specify the name of policy, the SVM to which it belongs, the FPolicy events to monitor, the FPolicy engine to which the generated notifications are sent and the policy priority. FPolicy policy configuration also allows to you to configure the file access behavior when the primary and secondary servers are down. Under such circumstances, if the "mandatory" field is set to true, file access is denied.

Each FPolicy policy is associated with a scope which allows you to restrain the scope of the policy to specified storage objects such as volume, shares and export or to a set of file extensions such as .txt, .jpeg. An FPolicy policy can be configured to send notifications, to the FPolicy server or for native file blocking which uses the file extension specified in the policy scope. An SVM can have multiple FPolicy policies which can be enabled or disabled independently of each other.

### Examples

#### Creating an FPolicy policy

Use the following API to create an FPolicy policy configuration. Note that the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"engine\": {
\"name\": \"engine1\" }, \"events\": [ { \"name\": \"cifs\" }, { \"name\":
\"nfs\" } ], \"mandatory\": true, \"name\": \"FPolicy_policy_0\",
\"passthrough_read\": true, \"privileged_user\": \"mydomain\\\\testuser\",
\"scope\": { \"exclude_export_policies\": [ \"export_pol1\" ],
\"exclude_extension\": [ \"txt\", \"png\" ], \"exclude_shares\": [ \"sh1\"
], \"exclude_volumes\": [ \"vol0\" ], \"include_export_policies\": [
\"export_pol10\" ], \"include_extension\": [ \"pdf\" ],
\"include_shares\": [ \"sh2\", \"sh3\" ], \"include_volumes\": [ \"vol1\",
\"vol2\" ] }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "FPolicy_policy_0",
```

```
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
      ],
    "engine": {
      "name": "engine1"
    },
    "scope": {
      "include_shares": [
        "sh2",
        "sh3"
      ],
      "exclude_shares": [
        "sh1"
      ],
      "include_volumes": [
        "vol1",
        "vol2"
      ],
      "exclude_volumes": [
        "vol0"
      ],
      "include_export_policies": [
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "privileged_user": "mydomain\\testuser",
    "passthrough_read": true
  }
]
}
```

**Creating and enable an FPolicy policy**

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"priority\":
1, \"engine\": { \"name\": \"engine1\" }, \"events\": [ { \"name\":
\"cifs\" }, { \"name\": \"nfs\" } ], \"mandatory\": true, \"name\":
\"FPolicy_policy_on\", \"passthrough_read\": false, \"scope\": {
\"exclude_export_policies\": [ \"export_pol1\" ], \"exclude_extension\": [
\"txt\", \"png\" ], \"exclude_shares\": [ \"sh1\" ], \"exclude_volumes\":
[ \"vol0\" ], \"include_export_policies\": [ \"export_pol10\" ],
\"include_extension\": [ \"pdf\" ], \"include_shares\": [ \"sh2\", \"sh3\"
], \"include_volumes\": [ \"vol1\", \"vol2\" ] }}"

# The response:
{
"num_records": 1,
"records": [
  {
   "name": "FPolicy_policy_0",
   "priority": 1,
   "events": [
     {
       "name": "cifs"
     },
     {
       "name": "nfs"
     }
   ],
   "engine": {
     "name": "engine1"
   },
   "scope": {
     "include_shares": [
       "sh2",
       "sh3"
     ],
     "exclude_shares": [
       "sh1"
     ],
```

```
      "include_volumes": [
        "vol1",
        "vol2"
      ],
      "exclude_volumes": [
        "vol0"
      ],
      "include_export_policies": [
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "privileged_user": "mydomain\\testuser",
    "passthrough_read": true
  }
  ]
}
```

#### Creating an FPolicy policy with the minimum required fields and a native engine

```
# The API:
POST /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"events\": [
{ \"name\": \"cifs\" }, { \"name\": \"nfs\" } ], \"name\":
\"pol_minimum_fields\", \"scope\": { \"include_volumes\": [ \"vol1\",
\"vol2\" ] }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "pol_minimum_fields",
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
    ],
    "scope": {
      "include_volumes": [
        "vol1",
        "vol2"
      ]
    }
  }
]
}
```

#### Retrieving all the FPolicy policy configurations for an SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/policies

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "pol0",
    "enabled": false,
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
    ],
    "engine": {
      "name": "engine1"
    },
    "scope": {
      "include_shares": [
        "sh2",
        "sh3"
      ],
      "exclude_shares": [
        "sh1"
      ],
      "include_volumes": [
        "vol1",
        "vol2"
      ],
      "exclude_volumes": [
        "vol0"
      ],
      "include_export_policies": [
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
```

```
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  },
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "FPolicy_policy_on",
    "enabled": true,
    "priority": 1,
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
    ],
    "engine": {
      "name": "engine1"
    },
     "scope": {
      "include_shares": [
        "sh2",
        "sh3"
      ],
      "exclude_shares": [
        "sh1"
      ],
      "include_volumes": [
        "vol1",
        "vol2"
      ],
      "exclude_volumes": [
        "vol0"
      ],
      "include_export_policies": [
```

```
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  },
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "cluster_pol",
    "enabled": false,
    "events": [
      {
        "name": "cluster"
      }
    ],
    "engine": {
      "name": "native"
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  },
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "pol_minimum_fields",
    "enabled": false,
    "events": [
      {
        "name": "cifs"
```

```
      },
      {
        "name": "nfs"
      }
    ],
    "engine": {
      "name": "native"
    },
    "scope": {
      "include_volumes": [
        "vol1",
        "vol2"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  }
],
"num_records": 4
}
```

#### Retrieving all of the FPolicy policy configurations for the FPolicy engine "engine1" for an SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-
0050568eeb34/policies?engine.name=engine1&fields=*&return_records=true&ret
urn_timeout=15" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "pol0",
```

```
    "enabled": false,
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
    ],
    "engine": {
      "name": "engine1"
    },
    "scope": {
      "include_export_policies": [
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  },
  {
    "svm": {
      "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "FPolicy_policy_on",
    "enabled": true,
    "priority": 1,
    "events": [
      {
        "name": "cifs"
      },
      {
        "name": "nfs"
      }
```

```
    ],
    "engine": {
      "name": "engine1"
    },
    "scope": {
      "include_shares": [
        "sh2",
        "sh3"
      ],
      "exclude_shares": [
        "sh1"
      ],
      "include_volumes": [
        "vol1",
        "vol2"
      ],
      "exclude_volumes": [
        "vol0"
      ],
      "include_export_policies": [
        "export_pol10"
      ],
      "exclude_export_policies": [
        "export_pol1"
      ],
      "include_extension": [
        "pdf"
      ],
      "exclude_extension": [
        "txt",
        "png"
      ]
    },
    "mandatory": true,
    "passthrough_read": false,
    "allow_privileged_access": false,
    "persistent_store": "ps1"
  }
],
"num_records": 2
}
```

#### Retrieving a particular FPolicy policy configuration for an SVM

```
# The API:
GET /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies/pol0" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
},
"name": "pol0",
"enabled": false,
"events": [
  {
    "name": "cifs"
  },
  {
    "name": "nfs"
  }
],
"engine": {
  "name": "engine1"
},
"scope": {
  "include_shares": [
    "sh2",
    "sh3"
  ],
  "exclude_shares": [
    "sh1"
  ],
  "include_volumes": [
    "vol1",
    "vol2"
  ],
  "exclude_volumes": [
    "vol0"
  ],
  "include_export_policies": [
    "export_pol10"
```

```
  ],
  "exclude_export_policies": [
    "export_pol1"
  ],
  "include_extension": [
     "pdf"
  ],
  "exclude_extension": [
    "txt",
    "png"
  ]
},
"mandatory": true,
"passthrough_read": false,
"allow_privileged_access": false,
"persistent_store": "ps1"
}
```

### Updating a particular FPolicy policy

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"engine\": { \"name\": \"native\"
}, \"events\": [ { \"name\": \"cifs\" } ], \"mandatory\": false,
\"scope\": { \"include_volumes\": [ \"*\" ] }}"
```

### Enabling a particular FPolicy policy

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"enabled\": true, \"priority\":
3}"
```

---

#### Disabling a particular FPolicy policy

---

```
# The API:
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"enabled\": true }"
```

---

### Retrieve an FPolicy configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}/policies`

**Introduced In:** 9.6

Retrieves the FPolicy policy configuration of an SVM. ONTAP allows the creation of a cluster level FPolicy policy that acts as a template for all the data SVMs belonging to the cluster. This cluster level FPolicy policy is also retrieved for the specified SVM.

#### Related ONTAP commands

- `fpolicy policy show`
- `fpolicy policy scope show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| allow_privileged_access | boolean | query | False | Filter by allow_privileged_access<br>• Introduced in: 9.13 |
| privileged_user | string | query | False | Filter by privileged_user<br>• Introduced in: 9.10 |
| enabled | boolean | query | False | Filter by enabled |
| engine.name | string | query | False | Filter by engine.name |
| name | string | query | False | Filter by name |
| priority | integer | query | False | Filter by priority<br>• Max value: 10<br>• Min value: 1 |
| passthrough_read | boolean | query | False | Filter by passthrough_read<br>• Introduced in: 9.10 |
| scope.object_monitoring_with_no_extension | boolean | query | False | Filter by scope.object_monitoring_with_no_extension<br>• Introduced in: 9.11 |
| scope.exclude_volumes | string | query | False | Filter by scope.exclude_volumes |
| scope.exclude_shares | string | query | False | Filter by scope.exclude_shares |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scope.include_shares | string | query | False | Filter by scope.include_shares |
| scope.exclude_extension | string | query | False | Filter by scope.exclude_extension |
| scope.include_volumes | string | query | False | Filter by scope.include_volumes |
| scope.check_extensions_on_directories | boolean | query | False | Filter by scope.check_extensions_on_directories<br>• Introduced in: 9.11 |
| scope.exclude_export_policies | string | query | False | Filter by scope.exclude_export_policies |
| scope.include_export_policies | string | query | False | Filter by scope.include_export_policies |
| scope.include_extension | string | query | False | Filter by scope.include_extension |
| mandatory | boolean | query | False | Filter by mandatory |
| events.name | string | query | False | Filter by events.name |
| persistent_store | string | query | False | Filter by persistent_store<br>• Introduced in: 9.14 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of Records |
| records | array[fpolicy_policy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "events": [
        "event_cifs",
        "event_open"
      ],
      "name": "fp_policy_1",
      "persistent_store": "ps1",
      "priority": 1,
      "privileged_user": "mydomain\\testuser",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
```

```
          "string"
        ],
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [
          "vol1",
          "vol_svm1"
        ]
      },
      "svm": {
        "uuid": "string"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

scope

| Name | Type | Description |
| --- | --- | --- |
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true, all objects with or without extensions are monitored. Default is false. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_policy

| Name | Type | Description |
| --- | --- | --- |
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes.<br>When this parameter is set to true,<br>FPolicy servers can access files on the cluster using a separate data channel with privileged access. |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period.<br>When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy.<br>Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |

| Name | Type | Description |
|---|---|---|
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an FPolicy configuration for an SVM**

POST `/protocols/fpolicy/{svm.uuid}/policies`

**Introduced In:** 9.6

Creates an FPolicy policy configuration for the specified SVM. To create an FPolicy policy, you must specify the policy scope and the FPolicy events to be monitored.

Important notes:

- A single policy can monitor multiple events.
- An FPolicy engine is an optional field whose default value is set to native. A native engine can be used to simply block the file access based on the file extensions specified in the policy scope.
- To enable a policy, the policy priority must be specified. If the priority is not specified, the policy is created but it is not enabled.
- The "mandatory" field, if set to true, blocks the file access when the primary or secondary FPolicy servers are down.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the FPolicy policy.
- `events` - Name of the events to monitor.
- `name` - Name of the FPolicy policy.
- `scope` - Scope of the policy. Can be limited to exports, volumes, shares or file extensions.
- `priority`- Priority of the policy (ranging from 1 to 10).

#### Default property values

- `mandatory` - *true*
- `engine` - *native*

#### Related ONTAP commands

- `fpolicy policy scope create`
- `fpolicy policy create`
- `fpolicy enable`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy. Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used<br>to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |

| Name | Type | Description |
| --- | --- | --- |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

**Example request**

```
{
  "allow_privileged_access": true,
  "enabled": true,
  "engine": {
    "name": "string"
  },
  "events": [
    "event_cifs",
    "event_open"
  ],
  "name": "fp_policy_1",
  "persistent_store": "ps1",
  "priority": 1,
  "privileged_user": "mydomain\\testuser",
  "scope": {
    "exclude_export_policies": [
      "string"
    ],
    "exclude_extension": [
      "string"
    ],
    "exclude_shares": [
      "string"
    ],
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": [
      "string"
    ],
    "include_extension": [
      "string"
    ],
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
      "vol_svm1"
    ]
  },
```

```
  "svm": {
    "uuid": "string"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of Records |
| records | array[fpolicy_policy] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "allow_privileged_access": true,
      "enabled": true,
      "engine": {
        "name": "string"
      },
      "events": [
        "event_cifs",
        "event_open"
      ],
      "name": "fp_policy_1",
      "persistent_store": "ps1",
      "priority": 1,
      "privileged_user": "mydomain\\testuser",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
          "string"
        ],
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [
          "vol1",
```

```
            "vol_svm1"
          ]
        },
        "svm": {
          "uuid": "string"
        }
      }
    ]
  }
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9764875 | An FPolicy event does not exist |
| 9764888 | An FPolicy engine does not exist |
| 9764898 | An FPolicy policy cannot be created without defining its scope |
| 9765027 | FPolicy creation is successful but it cannot be enabled as the priority is already in use by another policy |
| 9765037 | FPolicy creation failed as passthrough-read cannot be enabled for policy without privileged user |
| 9765038 | Passthrough-read policies are not supported with asynchronous external engine |
| 9765056 | The specified Persistent Store does not exist |
| 9765059 | Persistent Store feature is not supported with native engine |
| 9765060 | Persistent Store feature is not supported with synchronous engine |
| 9765061 | Persistent Store feature is not supported with mandatory screening |

| Error Code | Description |
| --- | --- |
| 9765065 | A valid privileged user name must be in the form "domain-name\user-name" |
| 9765066 | The privileged user contains characters that are not allowed |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
|---|---|---|
| name | string | |

scope

| Name | Type | Description |
|---|---|---|
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true,<br>all objects with or without extensions are monitored. Default is false. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_policy

| Name | Type | Description |
| --- | --- | --- |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |

| Name | Type | Description |
| --- | --- | --- |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers registered for the policy. Passthrough-read is a way to read data for offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an FPolicy configuration for an SVM

DELETE `/protocols/fpolicy/{svm.uuid}/policies/{name}`

**Introduced In:** 9.6

Deletes a particular FPolicy policy configuration for a specified SVM. To delete a policy, you must first disable the policy.

#### Related ONTAP commands

- `fpolicy policy scope delete`
- `fpolicy policy delete`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 9764900 | Deletion of a cluster level FPolicy policy is not supported |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FPolicy configuration for an SVM

GET `/protocols/fpolicy/{svm.uuid}/policies/{name}`

**Introduced In:** 9.6

Retrieves a particular FPolicy policy configuration for a specified SVM. Cluster-level FPolicy policy configuration details cannot be retrieved for a data SVM.

#### Related ONTAP commands

- `fpolicy policy show`
- `fpolicy policy scope show`
- `fpolicy show`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes. When this parameter is set to true, FPolicy servers can access files on the cluster using a separate data channel with privileged access. |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| name | string | Specifies the name of the policy. |

| Name | Type | Description |
| --- | --- | --- |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy. Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used<br>to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

**Example response**


```
{
  "engine": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string"
  },
  "events": [
    "event_cifs",
    "event_open"
  ],
  "name": "fp_policy_1",
  "persistent_store": "ps1",
  "priority": 1,
  "privileged_user": "mydomain\\testuser",
  "scope": {
    "exclude_export_policies": [
      "string"
    ],
    "exclude_extension": [
      "string"
    ],
    "exclude_shares": [
      "string"
    ],
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": [
      "string"
    ],
    "include_extension": [
      "string"
    ],
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
```

```
        "vol_svm1"
      ]
    },
    "svm": {
      "uuid": "string"
    }
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

scope

| Name | Type | Description |
|---|---|---|
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |

| Name | Type | Description |
| --- | --- | --- |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true,<br>all objects with or without extensions are monitored. Default is false. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an FPolicy configuration for an SVM

PATCH `/protocols/fpolicy/{svm.uuid}/policies/{name}`

**Introduced In:** 9.6

Updates a particular FPolicy policy configuration for a specified SVM. PATCH can be used to enable or disable the policy. When enabling a policy, you must specify the policy priority. The policy priority of the policy is not required when disabling the policy. If the policy is enabled, the FPolicy policy cannot be modified.

#### Related ONTAP commands

- `fpolicy policy modify`
- `fpolicy policy scope modify`
- `fpolicy enable`
- `fpolicy disable`

#### Learn more

- DOC /protocols/fpolicy/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes. When this parameter is set to true, FPolicy servers can access files on the cluster using a separate data channel with privileged access. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers<br>registered for the policy. Passthrough-read is a way to read data for<br>offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used<br>to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

### Example request

```
{
  "engine": {
    "name": "string"
  },
  "events": [
    "event_cifs",
    "event_open"
  ],
  "persistent_store": "ps1",
  "priority": 1,
  "privileged_user": "mydomain\\testuser",
  "scope": {
    "exclude_export_policies": [
      "string"
    ],
    "exclude_extension": [
      "string"
    ],
    "exclude_shares": [
      "string"
    ],
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": [
      "string"
    ],
    "include_extension": [
      "string"
    ],
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
      "vol_svm1"
    ]
  },
  "svm": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 9764875 | An FPolicy event does not exist |
| 9764888 | An FPolicy engine does not exist |
| 9765026 | The priority must be specified when enabling the FPolicy policy |
| 9765025 | Cannot disable an FPolicy policy when the priority is specified |
| 9764899 | Cannot modify an enabled FPolicy policy |
| 9764899 | Deletion of a cluster policy is not supported |
| 9764907 | An FPolicy policy is already enabled |
| 9764908 | An FPolicy policy is already disabled |
| 9765029 | An FPolicy was modified but disable/enable failed as the policy is already disabled/enabled |
| 9765036 | Cannot modify an FPolicy policy as passthrough-read policies are not supported without privileged user |
| 9765038 | Passthrough-read policies are not supported with an external engine of type "asynchronous" |
| 9765039 | Passthrough-read policies are not supported with native engine |
| 9765056 | The specified Persistent Store does not exist |
| 9765062 | Policy with Persistent Store does not support mandatory screening |
| 9765065 | A valid privileged user name must be in the form "domain-name\user-name" |
| 9765066 | The privileged user contains characters that are not allowed |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

fpolicy_engine_reference

FPolicy external engine

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FPolicy external engine. |

fpolicy_event_reference

FPolicy events

| Name | Type | Description |
|---|---|---|
| name | string | |

scope

| Name | Type | Description |
|---|---|---|
| check_extensions_on_directories | boolean | Specifies whether the file name extension checks also apply to directory objects. If this parameter is set to true, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to false, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match. Default is false. |
| exclude_export_policies | array[string] | |
| exclude_extension | array[string] | |
| exclude_shares | array[string] | |
| exclude_volumes | array[string] | |
| include_export_policies | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| include_extension | array[string] | |
| include_shares | array[string] | |
| include_volumes | array[string] | |
| object_monitoring_with_no_extension | boolean | Specifies whether the extension checks also apply to objects with no extension. If this parameter is set to true,<br>all objects with or without extensions are monitored. Default is false. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | SVM UUID |

fpolicy_policy

| Name | Type | Description |
| --- | --- | --- |
| allow_privileged_access | boolean | Specifies whether privileged access is required for FPolicy servers.<br>Privileged access is used when the FPolicy server requires direct access to the cluster nodes.<br>When this parameter is set to true,<br>FPolicy servers can access files on the cluster using a separate data channel with privileged access. |
| enabled | boolean | Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled. |
| engine | fpolicy_engine_reference | FPolicy external engine |
| events | array[fpolicy_event_reference] | |

| Name | Type | Description |
| --- | --- | --- |
| mandatory | boolean | Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances. |
| passthrough_read | boolean | Specifies whether passthrough-read should be allowed for FPolicy servers registered for the policy. Passthrough-read is a way to read data for offline files without restoring the files to primary storage. Offline files are files that have been moved to secondary storage. |
| persistent_store | string | Specifies the persistent storage name. This can then be used to enable persistent mode for FPolicy events. |
| priority | integer | Specifies the priority that is assigned to this policy. |
| privileged_user | string | Specifies the privileged user name for accessing files on the cluster using a separate data channel with privileged access. The input for this field should be in "domain\username" format. |
| scope | scope | |
| svm | svm | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View lock details

### View lock details

#### Overview

A lock is a synchronization mechanism for enforcing limits on concurrent access to files where many clients can be
accessing the same file at the same time. Locks can be viewed or broken according to a wide range of query fields
that are presented in the lock information definition.

It is recommended that you provide as many fields as possible to optimize query processing.

#### Examples

### Retrieving locks with all fields for all SVMs

```
# The API:
GET /protocols/locks

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/locks?return_records=true&return_timeout=15&fields=*" -H
"accept: application/json"

# The response:
{
"records": [
```

```
  {
    "svm": {
      "uuid": "5b4b6847-9ae4-11eb-8597-0050568ec154",
      "name": "vs1"
    },
    "volume": {
      "uuid": "429dcc79-9af2-11eb-b313-0050568ec154",
      "name": "test_vol1"
    },
    "interface": {
      "name": "vs1.data",
      "uuid": "6bf26e25-9ae4-11eb-8597-0050568ec154",
      "ip": {
        "address": "10.140.115.95"
      }
    },
    "path": "/test_vol1/Demo406.TXT",
    "uuid": "bf03d8b4-e145-498a-902d-b9fe5d546d18",
    "constituent": false,
    "protocol": "cifs",
    "client_address": "10.74.7.22",
    "owner_id": "feff-
020600000002000000020400000000000000000000000ffff0a4a07161053010001000000
030000006c00000000000000",
    "share_lock": {
      "mode": "read_write_deny_write_delete",
      "soft": false
    },
    "smb": {
      "connect_state": "connected",
      "open_group_id":
"71756e2325a7eb11843f005056a4731c101068450bcdffff1c2c000000000000",
      "open_type": "durable"
    },
    "state": "granted",
    "type": "share_level",
    "node": {
      "name": "user-vsim3",
      "uuid": "1f29b875-9ae3-11eb-8597-0050568ec154"
    }
  },
  {
    "svm": {
      "uuid": "5b4b6847-9ae4-11eb-8597-0050568ec154",
      "name": "vs1"
    },
```

```
    "volume": {
      "uuid": "429dcc79-9af2-11eb-b313-0050568ec154",
      "name": "test_vol1"
    },
    "interface": {
      "name": "vs1.data",
      "uuid": "6bf26e25-9ae4-11eb-8597-0050568ec154",
      "ip": {
        "address": "10.140.115.95"
      }
    },
    "path": "/test_vol1/Demo406.TXT",
    "uuid": "0c33d18d-dcbc-492a-81b3-4c5740c46172",
    "constituent": false,
    "protocol": "cifs",
    "client_address": "10.74.7.22",
    "oplock_level": "batch",
    "owner_id": "feff-
020600000002000000020400000000000000000000000ffff0a4a07161053010001000000
030000006c00000000000000",
    "smb": {
      "connect_state": "connected",
      "open_group_id":
"71756e2325a7eb11843f005056a4731c101068450bcdffff1c2c000000000000"
    },
    "state": "granted",
    "type": "op_lock",
    "node": {
      "name": "user-vsim3",
      "uuid": "1f29b875-9ae3-11eb-8597-0050568ec154"
    }
  }
],
"num_records": 2
}
```


### Retrieving locks of a specific volume

```
# The API:
GET /protocols/locks

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/locks?volume.uuid=429dcc79-
9af2-11eb-b313-
0050568ec154&return_records=true&return_timeout=15&fields=*" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "5b4b6847-9ae4-11eb-8597-0050568ec154",
      "name": "vs1"
    },
    "volume": {
      "uuid": "429dcc79-9af2-11eb-b313-0050568ec154",
      "name": "test_vol1"
    },
    "interface": {
      "name": "vs1.data",
      "uuid": "6bf26e25-9ae4-11eb-8597-0050568ec154"
    },
    "path": "/test_vol1/Demo408.TXT",
    "uuid": "be1cdf00-37f1-4477-b6d0-bba8c4fa8c67"
  },
  {
    "svm": {
      "uuid": "5b4b6847-9ae4-11eb-8597-0050568ec154",
      "name": "vs1"
    },
    "volume": {
      "uuid": "429dcc79-9af2-11eb-b313-0050568ec154",
      "name": "test_vol1"
    },
    "interface": {
      "name": "vs1.data",
      "uuid": "6bf26e25-9ae4-11eb-8597-0050568ec154"
    },
    "path": "/test_vol1/Demo408.TXT",
    "uuid": "393cc06e-8b37-4f49-b09a-74d1eef79368"
  }
],
"num_records": 2
}
```

### Retrieving the lock for a specific UUID

```
# The API:
GET /protocols/locks/{uuid}

# The call:
curl  -X GET "https://<mgmt-ip>/api/protocols/locks/be1cdf00-37f1-4477-
b6d0-bba8c4fa8c67" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "5b4b6847-9ae4-11eb-8597-0050568ec154",
  "name": "vs1"
},
"volume": {
  "uuid": "429dcc79-9af2-11eb-b313-0050568ec154",
  "name": "test_vol1"
},
"interface": {
  "name": "vs1.data",
  "uuid": "6bf26e25-9ae4-11eb-8597-0050568ec154",
  "ip": {
    "address": "10.140.115.95"
  }
},
"path": "/test_vol1/Demo406.TXT",
"uuid": "be1cdf00-37f1-4477-b6d0-bba8c4fa8c67",
"constituent": false,
"protocol": "cifs",
"client_address": "10.74.7.22",
"owner_id": "feff-
020600000002000000020400000000000000000000000ffff0a4a07161053010001000000
030000006c00000000000000",
"share_lock": {
  "mode": "read_write_deny_write_delete",
  "soft": false
},
"smb": {
  "connect_state": "connected",
  "open_group_id":
"71756e2325a7eb11843f005056a4731c101068450bcdffff1c2c000000000000",
  "open_type": "durable"
},
```

```
"state": "granted",
"type": "share_level",
"node": {
  "name": "user-vsim3",
  "uuid": "1f29b875-9ae3-11eb-8597-0050568ec154"
}
}
```

#### Deleting the lock for a specific UUID

```
# The API:
DELETE /protocols/locks/{uuid}

# The call:
curl  -X DELETE "https://<mgmt-ip>/api/protocols/locks/be1cdf00-37f1-4477-
b6d0-bba8c4fa8c67" -H "accept: application/json"
```

#### Deleting all locks for a specific protocol

```
# The API:
DELETE /protocols/locks/{uuid}

# The call:
curl  -X DELETE "https://<mgmt-ip>/api/protocols/locks?protocol=cifs" -H
"accept: application/json"
```

### Retrieve lock details

GET `/protocols/locks`

**Introduced In:** 9.10

Retrieves locks details.

#### Related ONTAP commands

- `vserver locks show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |
| constituent | boolean | query | False | Filter by constituent |
| uuid | string | query | False | Filter by uuid |
| protocol | string | query | False | Filter by protocol |
| smb.connect_state | string | query | False | Filter by smb.connect_state |
| smb.open_type | string | query | False | Filter by smb.open_type |
| smb.open_group_id | string | query | False | Filter by smb.open_group_id |
| share_lock.mode | string | query | False | Filter by share_lock.mode |
| share_lock.soft | boolean | query | False | Filter by share_lock.soft |
| byte_lock.super | boolean | query | False | Filter by byte_lock.super |
| byte_lock.mandatory | boolean | query | False | Filter by byte_lock.mandatory |
| byte_lock.offset | integer | query | False | Filter by byte_lock.offset |
| byte_lock.soft | boolean | query | False | Filter by byte_lock.soft |
| byte_lock.exclusive | boolean | query | False | Filter by byte_lock.exclusive |
| byte_lock.length | integer | query | False | Filter by byte_lock.length |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interface.ip.address | string | query | False | Filter by interface.ip.address |
| interface.name | string | query | False | Filter by interface.name |
| interface.uuid | string | query | False | Filter by interface.uuid |
| state | string | query | False | Filter by state |
| client_address | string | query | False | Filter by client_address |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| type | string | query | False | Filter by type |
| path | string | query | False | Filter by path |
| delegation | string | query | False | Filter by delegation |
| oplock_level | string | query | False | Filter by oplock_level |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| owner_id | string | query | False | Filter by owner_id |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records. |
| records | array[client_lock] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "byte_lock": {
        "length": "10",
        "offset": "100"
      },
      "client_address": "0.0.0.0",
      "delegation": "string",
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "oplock_level": "string",
      "owner_id": "string",
      "path": "string",
      "protocol": "string",
      "share_lock": {
```

```
          "mode": "string"
        },
        "smb": {
          "connect_state": "string",
          "open_group_id": "string",
          "open_type": "string"
        },
        "state": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "type": "string",
        "uuid": "ceeac1b4-8646-4c76-a054-1c96e87594aa",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Unexpected error_response
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

byte_lock

| Name | Type | Description |
|---|---|---|
| exclusive | boolean | Indicates whether it is an exclusive bytelock. |
| length | integer | Length of the bytelock starting from the offset. |
| mandatory | boolean | Indicates whether or not the bytelock is mandatory. |
| offset | integer | Starting offset for a bytelock. |
| soft | boolean | Indicates whether it is a soft bytelock. |
| super | boolean | Indicates whether it is a super bytelock. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

share_lock

| Name | Type | Description |
|---|---|---|
| mode | string | Types of share lock modes. |
| soft | boolean | Indicates whether it is a soft share lock. |

smb

| Name | Type | Description |
|---|---|---|
| connect_state | string | SMB connection state. |
| open_group_id | string | SMB open group ID. |
| open_type | string | SMB open type. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

client_lock

This object represents locks on a volume.

| Name | Type | Description |
| --- | --- | --- |
| byte_lock | byte_lock | |
| client_address | string | IP address of the client holding the lock. |
| constituent | boolean | Indicate if volume is constituent or not. |
| delegation | string | Type of delegation. |

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| node | node | |
| oplock_level | string | The oplock level determines which operations the client may cache locally. |
| owner_id | string | Owner ID. |
| path | string | Object path |
| protocol | string | Type of lock protocol. |
| share_lock | share_lock | |
| smb | smb | |
| state | string | State of lock. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of lock. |
| uuid | string | Lock UUID |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete locks for a parameter

DELETE `/protocols/locks/{uuid}`

**Introduced In:** 9.10

Deletes locks of given parameter.

#### Related ONTAP commands

- `vserver locks break`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Lock ID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the lock for a specific UUID

GET `/protocols/locks/{uuid}`

**Introduced In:** 9.10

Retrieves the lock for a specific UUID.

#### Related ONTAP commands

- `vserver locks show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Lock ID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| byte_lock | byte_lock | |
| client_address | string | IP address of the client holding the lock. |
| constituent | boolean | Indicate if volume is constituent or not. |
| delegation | string | Type of delegation. |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| node | node | |
| oplock_level | string | The oplock level determines which operations the client may cache locally. |
| owner_id | string | Owner ID. |
| path | string | Object path |
| protocol | string | Type of lock protocol. |
| share_lock | share_lock | |

| Name | Type | Description |
| --- | --- | --- |
| smb | smb | |
| state | string | State of lock. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of lock. |
| uuid | string | Lock UUID |
| volume | volume | |

**Example response**


```
{
  "byte_lock": {
    "length": "10",
    "offset": "100"
  },
  "client_address": "0.0.0.0",
  "delegation": "string",
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "oplock_level": "string",
  "owner_id": "string",
  "path": "string",
  "protocol": "string",
  "share_lock": {
    "mode": "string"
  },
  "smb": {
    "connect_state": "string",
    "open_group_id": "string",
    "open_type": "string"
  },
  "state": "string",
  "svm": {
    "_links": {
      "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "type": "string",
    "uuid": "ceeac1b4-8646-4c76-a054-1c96e87594aa",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

byte_lock

| Name | Type | Description |
|---|---|---|
| exclusive | boolean | Indicates whether it is an exclusive bytelock. |
| length | integer | Length of the bytelock starting from the offset. |
| mandatory | boolean | Indicates whether or not the bytelock is mandatory. |
| offset | integer | Starting offset for a bytelock. |
| soft | boolean | Indicates whether it is a soft bytelock. |
| super | boolean | Indicates whether it is a super bytelock. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

share_lock

| Name | Type | Description |
| --- | --- | --- |
| mode | string | Types of share lock modes. |
| soft | boolean | Indicates whether it is a soft share lock. |

smb

| Name | Type | Description |
| --- | --- | --- |
| connect_state | string | SMB connection state. |
| open_group_id | string | SMB open group ID. |
| open_type | string | SMB open type. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## View NFS client information

### View NFS client information

#### Overview

ONTAP connected clients map functionality is mainly used by the System Manager to display client information.

The following are details of the fields retrieved for the Connected Clients MAP GET API:

node.name: Node name hosting this record, essentially the node hosting the "server_ip".
node.uuid: Node UUID hosting this record, essentially the node hosting the "server_ip".
svm.name: SVM name to which the "server_ip" belongs to.
svm.uuid: SVM UUID to which the "server_ip" belongs to.
server_ip: All clients that are connected to this interface are displayed in rows.
client_ips: List of client IP addresses connected to the interface.

#### Example

### Retrieving connected client Map information

```
# The API:
GET /protocols/nfs/connected-client-maps

# The call:
curl -X GET "https://<cluster-mgmt-ip>/api/protocols/nfs/connected-client-
maps?return_timeout=15&return_records=true" -H "accept: application/json"

# The response:
{
"records": [
 {
   "svm": {
     "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
     "name": "vs1"
   },
   "node": {
     "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
     "name": "vsim1"
   },
   "server_ip": "10.140.72.214",
   "client_ips": [
     "127.0.0.1"
   ]
 }],
 "num_records": 1
}
```

**Retrieve NFS client information**

GET `/protocols/nfs/connected-client-maps`

**Introduced In:** 9.11

Retrieves NFS clients information.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| client_ips | string | query | False | Filter by client_ips |
| server_ip | string | query | False | Filter by server_ip |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[nfs_clients_map] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_ips": [
        "127.0.0.1"
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "server_ip": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nfs_clients_map

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| client_ips | array[string] | Specifies the IP address of the client. |
| node | node | |
| server_ip | string | Specifies the IP address of the server. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update NFS client settings

### View and update NFS client settings

#### Overview

This cluster-wide API is used to set the maximum cache idle time (client_retention_interval) for the connected-clients cache. If a
client connected to NFS server is idle for longer than than the maximum cache idle time, the entry will be removed. The update_interval
value will change when the client_retention_interval is changed. The update interval represents the interval between the cleaning
happens. If the value of client_retention_interval is set to 60hrs the connected client entry will stay there for 60 hours
and after that it will get removed. If the value of update_interval is 8 hours then the cache will be refreshed

once every 8 hours.

#### Example

### Retrieves connected-client cache settings information

```
# The API:
GET /api/protocols/nfs/connected-client-settings

# The call:
curl -X GET "https://<cluster-mgmt-ip>/api/protocols/nfs/connected-client-
settings?return_timeout=15&return_records=true" -H "accept:
application/json"

# The response:
{
"client_retention_interval": "P7D",
"update_interval": "PT8H"
}
```

### Updating connected-client cache settings

```
# The API:
PATCH /api/protocols/nfs/connected-client-settings

# The call:
curl -d "@test_cache_settings_patch.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/connected-client-settings"
test_cache_settings_patch.txt(body):
{
"client_retention_interval": "P7D"
}
```

### Retrieve the NFS connected-client cache settings for a cluster

GET `/protocols/nfs/connected-client-settings`

**Introduced In:** 9.12

Retrieves the NFS connected-client cache settings of the cluster.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| client_retention_interval | string | query | False | Filter by client_retention_interval |
| update_interval | string | query | False | Filter by update_interval |
| enable_nfs_clients_deletion | boolean | query | False | Filter by enable_nfs_clients_deletion |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| client_retention_interval | string | The lifetime range of the connected-clients cache. Only intervals in multiples of 12 hours or its equivalent in days, minutes or seconds are supported. The minimum is 12 hours and the maximum is 168 hours or 7 days. |
| enable_nfs_clients_deletion | boolean | Specifies whether or not NFS Clients deletion is enabled for the connected-clients cache. When set to "true", connected-clients entries are deleted when a connection is closed. |
| update_interval | string | The time interval between refreshing the connected-clients cache. The minimum is 1 hour and the maximum is 8 hours. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "client_retention_interval": "string",
  "update_interval": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update NFS connected-client cache settings properties

PATCH `/protocols/nfs/connected-client-settings`

**Introduced In:** 9.12

Updates the properties of the NFS connected-client cache settings.

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| client_retention_interval | string | The lifetime range of the connected-clients cache. Only intervals in multiples of 12 hours or its equivalent in days, minutes or seconds are supported. The minimum is 12 hours and the maximum is 168 hours or 7 days. |
| enable_nfs_clients_deletion | boolean | Specifies whether or not NFS Clients deletion is enabled for the connected-clients cache. When set to "true", connected-clients entries are deleted when a connection is closed. |
| update_interval | string | The time interval between refreshing the connected-clients cache. The minimum is 1 hour and the maximum is 8 hours. |

**Example request**

```
{
  "client_retention_interval": "string",
  "update_interval": "string"
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 3277172 | The value for "client_retention_interval" must be between 12 and 168 hours or its equivalent in days, minutes or seconds. |

| Error Code | Description |
| --- | --- |
| 3277170 | The "client_retention_interval" must be set to a value which is a multiple of 12 hours or to its equivalent in days, minutes or seconds. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

nfs_clients_cache

| Name | Type | Description |
|---|---|---|
| client_retention_interval | string | The lifetime range of the connected-clients cache. Only intervals in multiples of 12 hours or its equivalent in days, minutes or seconds are supported. The minimum is 12 hours and the maximum is 168 hours or 7 days. |
| enable_nfs_clients_deletion | boolean | Specifies whether or not NFS Clients deletion is enabled for the connected-clients cache. When set to "true", connected-clients entries are deleted when a connection is closed. |
| update_interval | string | The time interval between refreshing the connected-clients cache. The minimum is 1 hour and the maximum is 8 hours. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View NFS connected clients

### View NFS connected clients

#### Overview

ONTAP connected clients show functionality is mainly used to provide a list of currently connected NFS clients. It also provides a potential list of other NFS clients that can be connected but are currently idle and a list of recently unmounted clients. If a client connected to the NFS server is idle for longer than the maximum cache idle time, then the entry will be removed. By default, the maximum cache idle time is 48 hours.

The following are details of the fields retrieved for the Connected Clients GET API:

node.name: The node name hosting this record; basically the node hosting the "server_ip".
node.uuid: The node UUID hosting this record; basically the node hosting the "server_ip".
svm.name: The svm name to which the "server_ip" belongs to.
svm.uuid: The svm uuid to which the "server_ip" belongs to.
server_ip: All clients that are connected to this interface are displayed in rows.
client_ip: The IP address of the client that is connected to the interface.
volume.name: The name of the volume the client is accessing.
volume.uuid: The UUID of the volume the client is accessing. This field is expensive field and will be fetched in advance privilege level.
protocol: The NFS protocol version over which client is accessing the volume.
export_policy.id: The export policy ID associated with the volume.
export_policy.name: The export policy name associated with the volume.
idle_duration: The time elapsed since the last request was sent by the client for this volume.
local_request_count: A counter that tracks requests that are sent to the volume with fast-path to local node.
remote_request_count: A counter that tracks requests that are sent to the volume with slow-path to remote node.
trunking_enabled: Flag that indicates the trunking status for the specified SVM connection. True indicates that the trunking feature is enabled while false indicates that the trunking feature is disabled.

#### Related CLI command

The CLI command for ONTAP connected clients is: "nfs connected-clients show"
cluster-1::*> nfs connected-clients show
Node: node1
Vserver: vs1
Data-Ip: xxx.xxx.xxx.xxx
Local Remote
Client-Ip Protocol Volume Policy Idle-Time Reqs Reqs Trunking
---------- ------ ------ ------ -------- ---- ---- ----

xxx.xxx.xxx.xxx nfs4.1 root_vol pol1 6s 9 0 false
xxx.xxx.xxx.xxx nfs3 vol1 pol1 56s 7 0 false
xxx.xxx.xxx.xxx nfs4.1 vol1 pol1 6s 7 0 false
3 entries were displayed.

#### Example

### Retrieves connected client information

```
# The API:
GET /protocols/nfs/connected-clients

# The call:
curl -X GET "https://<cluster-mgmt-ip>/api/protocols/nfs/connected-
clients?return_timeout=15&return_records=true" -H "accept:
application/json"

# The response:
{
"records": [
{
   "svm": {
     "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
     "name": "vs1"
   },
   "node": {
     "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
     "name": "vsim1"
   },
   "server_ip": "10.140.72.214",
   "client_ip": "10.140.137.57",
   "volume": {
     "name": "rvol1",
     "uuid": "c6bbc6f2-b8d0-11e9-9ad1-0050568e8480"
   },
   "protocol": "nfs4"
 },
 {
   "svm": {
     "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
     "name": "vs1"
   },
   "node": {
     "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
     "name": "vsim1"
   },
   "server_ip": "10.140.72.214",
```

```
    "client_ip": "10.140.137.57",
    "volume": {
      "name": "vol1",
      "uuid": "d28d1999-b8d0-11e9-9ad1-0050568e8480"
    },
    "protocol": "nfs3"
  },
  {
    "svm": {
      "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
      "name": "vs1"
    },
    "node": {
      "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
      "name": "vsim1"
    },
    "server_ip": "10.140.72.214",
    "client_ip": "10.140.137.57",
    "volume": {
      "name": "vol1",
      "uuid": "d28d1999-b8d0-11e9-9ad1-0050568e8480"
    },
    "protocol": "nfs4"
  }],
  "num_records": 3
}
```


**Retrieve the NFS configuration for an SVM**

GET `/protocols/nfs/connected-clients`

**Introduced In:** 9.7

Retrieves the NFS configuration of SVMs.

#### Expensive properties

export_policy.id is expensive field. It is not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `export_policy.id`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.name | string | query | False | Filter by volume.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| volume.uuid | string | query | False | Filter by volume.uuid |
| client_ip | string | query | False | Filter by client_ip |
| trunking_enabled | boolean | query | False | Filter by trunking_enabled<br>• Introduced in: 9.12 |
| protocol | string | query | False | Filter by protocol |
| export_policy.name | string | query | False | Filter by export_policy.name<br>• Introduced in: 9.9 |
| export_policy.id | integer | query | False | Filter by export_policy.id<br>• Introduced in: 9.9 |
| idle_duration | string | query | False | Filter by idle_duration |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| server_ip | string | query | False | Filter by server_ip |
| remote_request_count | integer | query | False | Filter by remote_request_count |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| local_request_count | integer | query | False | Filter by local_request_count |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[nfs_clients] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_ip": "string",
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 100,
        "name": "default"
      },
      "idle_duration": "P4DT84H30M5S",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "protocol": "nfs3",
      "server_ip": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917887 | The first character of the volume name must be a letter or underscore. |
| 917888 | Invalid character in volume name. It can have '_' and alphanumeric characters. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

export_policy

Export Policy

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| id | integer | |
| name | string | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

nfs_clients

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| client_ip | string | Specifies IP address of the client. |
| export_policy | export_policy | Export Policy |
| idle_duration | string | Specifies an ISO-8601 format of date and time to retrieve the idle time duration in hours, minutes, and seconds format. |
| local_request_count | integer | A counter that tracks requests that are sent to the volume with fast-path to local node. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The NFS protocol version over which client is accessing the volume. The following values are supported:<br>• nfs - All NFS versions are considered<br>• nfs3 - NFS version 3 protocol<br>• nfs4 - NFS version 4 protocol<br>• nfs4.1 - NFS version 4 minor version 1 protocol<br>• nfs4.2 - NFS version 4 minor version 2 protocol |
| remote_request_count | integer | A counter that tracks requests that are sent to the volume with slow-path to remote node. |
| server_ip | string | Specifies the IP address of the server. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| trunking_enabled | boolean | Flag that indicates the trunking status for the specified SVM connection. True indicates that the trunking feature is enabled while false indicates that the trunking feature is disabled. |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NFS export policies

### Manage NFS export policies

#### Export Policies

### 1) Retrieve the export policy details

### NOTE: This is used to read in the whole policy.

```
# The API:
GET /api/protocols/nfs/export-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-policies"
```

### 2) Create an export policy for an SVM

```
# The API:
POST /api/protocols/nfs/export-policies

# The call:
curl -d "@test_post_policy_single_rule.txt" -X POST "https://<mgmt-
ip>/api/protocols/nfs/export-policies"
test_post_policy_single_rule.txt(body):
{
"name": "P1",
"rules":[
  {
    "clients": [
      {
```

```
        "match": "host1"
      }
    ],
    "ro_rule": [
      "krb5"
    ],
    "rw_rule": [
      "ntlm"
    ],
    "anonymous_user": "anon1",
    "chown_mode": "restricted",
    "allow_suid": true
  },
  {
    "clients": [
      {
        "match": "host2"
      }
    ],
    "ro_rule": [
      "sys"
    ],
    "rw_rule": [
      "ntlm"
    ],
    "superuser": [
      "any"
    ],
    "allow_device_creation": true,
    "ntfs_unix_security": "fail"
  }
]
}
```


---

**3) Update an export policy for an SVM**

**NOTE: This is used to either update the policy name or update all of the rules at once in an existing export policy.**

---

```
# The API:
PATCH /api/protocols/nfs/export-policies/{policy.id}

# The call:
curl -d "@test_patch_policy.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934594"
test_patch_policy.txt(body):
{
"name": "S1",
"rules":[
  {
    "clients": [
      {
        "match": "host4"
      }
    ],
    "ro_rule": [
      "krb5"
    ],
    "rw_rule": [
      "ntlm"
    ]
  }
]
}
```

#### 4) Delete an export policy for an SVM

```
# The API:
DELETE /api/protocols/nfs/export-policies/{policy.id}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934594"
```

#### Export Rules

**1) Retrieve the export policy rule details for an export policy**

```
# The API:
GET /api/protocols/nfs/export-policies/{policy.id}/rules

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934595/rules"
```

**2) Create an export policy rule for an export policy**

**NOTE: This is used to add a single new rule to an existing export policy.**

```
# The API:
POST /api/protocols/nfs/export-policies/{policy.id}/rules

# The call:
curl -d "@test_patch_export_rule.txt" -X POST "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934595/rules"
test_patch_export_rule.txt(body):
{
"clients": [
  {
    "match": "host2"
  }
],
"ro_rule": [
  "sys"
],
"rw_rule": [
  "ntlm"
]
}
```

**3) Update an export policy rule for an export policy**

**NOTE: This is used to edit an existing single rule in an existing export policy.**

```
# The API:
PATCH /api/protocols/nfs/export-policies/{policy.id}/rules/{index}

# The call:
curl -d "@test_patch_export_rule.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934595/rules/5?new_index=10"
test_patch_export_rule.txt(body):
{
"clients": [
  {
    "match": "host4"
  }
],
"ro_rule": [
  "sys"
],
"rw_rule": [
  "krb5"
]
}
```

---

**4) Delete an export policy rule for an export policy**

---

```
# The API:
DELETE /api/protocols/nfs/export-policies/{policy.id}/rules/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934595/rules/15"
```

---

#### Export Clients

**1) Retrieve the export client matches of an export policy rule**

---

```
# The API:
GET /api/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934593/rules/2/clients"
```

---

**2) Add an export client match to an export policy rule**

---

```
# The API:
POST /api/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

# The call:
curl -d "@add_client_match.txt" -X POST "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934593/rules/1/clients"
add_client_match.txt(body):
{
"match" : "host4"
}
```

---

**3) Delete an export client match from an export policy rule**

---

```
# The API:
DELETE /api/protocols/nfs/export-
policies/{policy.id}/rules/{index}/clients/{match}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934593/rules/1/clients/host1,host2"
```

---

**Retrieve export policies**

GET `/protocols/nfs/export-policies`

**Introduced In:** 9.6

---

Retrieves export policies.

#### Related ONTAP commands

- `vserver export-policy show`
- `vserver export-policy rule show`

#### Learn more

- DOC /protocols/nfs/export-policies

> `are_rules_truncated` is set to 'true' if the total number of export-policy rules is greater than 500. If set to 'true', rules after 500 will be truncated. Use GET /api/protocols/nfs/export-policies/{id}/rules to retrieve all rules for truncated policies.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | query | False | Filter by id |
| are_rules_truncated | boolean | query | False | Filter by are_rules_truncated<br>• Introduced in: 9.18 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name<br>• maxLength: 256 |
| rules.index | integer | query | False | Filter by rules.index |
| rules.clients.match | string | query | False | Filter by rules.clients.match |
| rules.chown_mode | string | query | False | Filter by rules.chown_mode<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rules.ntfs_unix_security | string | query | False | Filter by rules.ntfs_unix_security<br>• Introduced in: 9.9 |
| rules.allow_suid | boolean | query | False | Filter by rules.allow_suid<br>• Introduced in: 9.9 |
| rules.superuser | string | query | False | Filter by rules.superuser |
| rules.ro_rule | string | query | False | Filter by rules.ro_rule |
| rules.protocols | string | query | False | Filter by rules.protocols |
| rules.anonymous_user | string | query | False | Filter by rules.anonymous_user |
| rules.allow_device_creation | boolean | query | False | Filter by rules.allow_device_creation<br>• Introduced in: 9.9 |
| rules.rw_rule | string | query | False | Filter by rules.rw_rule |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of export policy records |
| records | array[export_policy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 0,
      "name": "string",
      "rules": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "anonymous_user": "string",
          "chown_mode": "string",
          "clients": [
            {
              "match": "0.0.0.0/0"
            }
          ],
          "ntfs_unix_security": "string",
          "protocols": [
            "string"
          ],
          "ro_rule": [
            "string"
          ],
          "rw_rule": [
            "string"
          ],
          "superuser": [
```

```
            "string"
          ]
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262197 | The value provided is invalid for the field |
| 2621462 | The specified SVM name does not exist |
| 2621519 | SVM name is invalid. The SVM name must begin with a letter or an underscore. If the SVM is of type "sync-source", the maximum supported length is 41. Otherwise, the maximum supported length is 47 |
| 2621643 | The specified SVM name is too long |
| 2621685 | SVM name length cannot be zero |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 6691623 | User is not authorized |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| next | href | |
| self | href | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |

| Name | Type | Description |
|---|---|---|
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an export policy

POST `/protocols/nfs/export-policies`

**Introduced In:** 9.6

Creates an export policy. An SVM can have any number of export policies to define rules for which clients can access data exported by the SVM. A policy with no rules prohibits access.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create an export policy.
- `name` - Name of the export policy.

#### Recommended optional properties

- `rules` - Rule(s) of an export policy. Used to create the export rule and populate the export policy with export rules in a single request.
- `rules[].index` - If you specify an index number that already matches a rule, the index number of the existing rule is incremented, as are the index numbers of all subsequent rules, either to the end of the list or to an open space in the list. If you do not specify an index number, the new rule is placed at the end of the policy's list.

#### Related ONTAP commands

- `vserver export-policy create`
- `vserver export-policy rule create`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "id": 0,
  "name": "string",
  "rules": [
    {
      "anonymous_user": "string",
      "chown_mode": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "ntfs_unix_security": "string",
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ]
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of export policy records |

| Name | Type | Description |
| --- | --- | --- |
| records | array[export_policy] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "id": 0,
      "name": "string",
      "rules": [
        {
          "anonymous_user": "string",
          "chown_mode": "string",
          "clients": [
            {
              "match": "0.0.0.0/0"
            }
          ],
          "ntfs_unix_security": "string",
          "protocols": [
            "string"
          ],
          "ro_rule": [
            "string"
          ],
          "rw_rule": [
            "string"
          ],
          "superuser": [
            "string"
          ]
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703952 | Invalid ruleset name provided. No spaces allowed in a ruleset name |
| 1703954 | Export policy does not exist |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704047 | The export policy name cannot be longer than 256 characters |
| 1704049 | Invalid clientmatch: clientmatch lists require an effective cluster version of Data ONTAP 9.0 or later. Upgrade all nodes to Data ONTAP 9.0 or above to use features that operate on lists of clientmatch strings in export-policy rules |
| 1704050 | Invalid clientmatch: clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |

| Error Code | Description |
| --- | --- |
| 1704055 | Export policies are only supported for data Vservers |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 1704071 | The export policy was not created because of duplicate export policy rules present in the request |
| 2621462 | The specified SVM name does not exist |
| 2621519 | SVM name is invalid. The SVM name must begin with a letter or an underscore. If the SVM is of type "sync-source", the maximum supported length is 41. Otherwise, the maximum supported length is 47 |
| 2621643 | The specified SVM name is too long |
| 2621685 | SVM name length cannot be zero |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 3277000 | Upgrade all nodes to ONTAP 9.0.0 or above to use krb5p as a security flavor in export-policy rules |
| 3277083 | User ID is not valid. Enter a value for User ID from 0 to 4294967295 |
| 3277149 | The "Anon" field cannot be an empty string |
| 6691623 | User is not authorized |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_policy

| Name | Type | Description |
|---|---|---|
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an export policy**

DELETE `/protocols/nfs/export-policies/{id}`

**Introduced In:** 9.6

Deletes an export policy.

#### Related ONTAP commands

- `vserver export-policy delete`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Export Policy ID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703944 | Failed to delete rule |
| 1703945 | Ruleset is in use by a volume. It cannot be deleted until all volumes that refer to it are first deleted |
| 1703946 | Cannot determine if the ruleset is in use by a volume. It cannot be deleted until all volumes that refer to it are first deleted |
| 1703947 | Cannot delete default ruleset. This ruleset will be deleted when the owning Vserver is deleted |
| 1703952 | Invalid ruleset name provided. No spaces are allowed in a ruleset name |
| 1703953 | This ruleset is in use by a qtree export policy. It cannot be deleted until all qtree policies that refer to it are first deleted |
| 1704047 | The export policy name cannot be longer than 256 characters |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an export policy

GET `/protocols/nfs/export-policies/{id}`

**Introduced In:** 9.6

Retrieves an export policy.

#### Related ONTAP commands

- `vserver export-policy show`
- `vserver export-policy rule show`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Export Policy ID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 0,
  "name": "string",
  "rules": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anonymous_user": "string",
      "chown_mode": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "ntfs_unix_security": "string",
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ]
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

export_clients

| Name | Type | Description |
|---|---|---|
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update export policy properties

PATCH `/protocols/nfs/export-policies/{id}`

**Introduced In:** 9.6

Updates the properties of an export policy to change an export policy name or replace all export policy rules.

#### Related ONTAP commands

- `vserver export-policy rename`
- `vserver export-policy rule delete`
- `vserver export-policy rule create`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Export Policy ID |

**Request Body**

| Name | Type | Description |
|---|---|---|
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "id": 0,
  "name": "string",
  "rules": [
    {
      "anonymous_user": "string",
      "chown_mode": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "ntfs_unix_security": "string",
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ]
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703950 | Failed to rename ruleset |
| 1703952 | Invalid ruleset name provided. No spaces are allowed in a ruleset name |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704047 | The export policy name cannot be longer than 256 characters |
| 1704050 | Invalid clientmatch: clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 3277000 | Upgrade all nodes to ONTAP 9.0.0 or above to use krb5p as a security flavor in export-policy rules |
| 3277083 | User ID is not valid. Enter a value for User ID from 0 to 4294967295 |
| 3277149 | The "Anon" field cannot be an empty string |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

export_rules

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_policy

| Name | Type | Description |
|---|---|---|
| are_rules_truncated | boolean | Indicates whether or not export-policy rules are truncated. |
| id | integer | Export Policy ID |
| name | string | Export Policy Name |
| rules | array[export_rules] | Rules of the Export Policy. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve export policy rules**

GET `/protocols/nfs/export-policies/{policy.id}/rules`

**Introduced In:** 9.6

Retrieves export policy rules.

#### Related ONTAP commands

- `vserver export-policy rule show`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| allow_device_creation | boolean | query | False | Filter by allow_device_creation<br>• Introduced in: 9.9 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.10 |
| svm.uuid | string | query | False | Filter by svm.uuid<br>• Introduced in: 9.10 |
| rw_rule | string | query | False | Filter by rw_rule |
| superuser | string | query | False | Filter by superuser |
| policy.name | string | query | False | Filter by policy.name<br>• Introduced in: 9.10 |
| protocols | string | query | False | Filter by protocols |
| ro_rule | string | query | False | Filter by ro_rule |
| anonymous_user | string | query | False | Filter by anonymous_user |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| allow_suid | boolean | query | False | Filter by allow_suid<br>• Introduced in: 9.9 |
| index | integer | query | False | Filter by index |
| clients.match | string | query | False | Filter by clients.match |
| chown_mode | string | query | False | Filter by chown_mode<br>• Introduced in: 9.9 |
| ntfs_unix_security | string | query | False | Filter by ntfs_unix_security<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of Export Rule records |
| records | array[export_rule] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anonymous_user": "string",
      "chown_mode": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "ntfs_unix_security": "string",
      "policy": {
        "name": "string"
      },
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262194 | The expression is missing a value |
| 262197 | The value provided is invalid for the field |
| 2621462 | The specified SVM name does not exist |
| 2621519 | SVM name is invalid. The SVM name must begin with a letter or an underscore. If the SVM is of type "sync-source", the maximum supported length is 41. Otherwise, the maximum supported length is 47 |
| 2621643 | The specified SVM name is too long |
| 2621685 | SVM name length cannot be zero |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 6691623 | User is not authorized |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| next | href | |
| self | href | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Export policy ID |
| name | string | Export policy name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_rule

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| policy | policy | |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an export policy rule

POST `/protocols/nfs/export-policies/{policy.id}/rules`

**Introduced In:** 9.6

Creates an export policy rule.

#### Required properties

- `policy.id` - Existing export policy for which to create an export rule.
- `clients.match` - List of clients (hostnames, ipaddresses, netgroups, domains) to which the export rule applies.
- `ro_rule` - Used to specify the security type for read-only access to volumes that use the export rule.
- `rw_rule` - Used to specify the security type for read-write access to volumes that use the export rule.

#### Optional property

- `index` - Used to specify the index number of the export rule that you want to create. If you specify an index number that already matches a rule, the index number of the existing rule is incremented, as are the index numbers of all subsequent rules, either to the end of the list or to an open space in the list. If you do not

specify an index number, the new rule is placed at the end of the policy's list.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `protocols` - *any*
- `anonymous_user` - *none*
- `superuser` - *any*
- `allow_device_creation` - *true*
- `ntfs_unix_security` - *fail*
- `chown_mode` - *restricted*
- `allow_suid` - *true*

#### Related ONTAP commands

- `vserver export-policy rule create`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| policy | policy | |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "anonymous_user": "string",
  "chown_mode": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ],
  "ntfs_unix_security": "string",
  "policy": {
    "name": "string"
  },
  "protocols": [
    "string"
  ],
  "ro_rule": [
    "string"
  ],
  "rw_rule": [
    "string"
  ],
  "superuser": [
    "string"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of Export Rule records |
| records | array[export_rule] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "anonymous_user": "string",
      "chown_mode": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "ntfs_unix_security": "string",
      "policy": {
        "name": "string"
      },
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703954 | Export policy does not exist |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704050 | Invalid clientmatch: clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 1704070 | Export policy rule already exists. The export policy rule was not created because an export policy rule with the same values already exists. |
| 3277000 | Upgrade all nodes to ONTAP 9.0.0 or above to use krb5p as a security flavor in export-policy rules |
| 3277083 | User ID is not valid. Enter a value for User ID from 0 to 4294967295 |
| 3277162 | The specified "index", 0, is invalid. Valid values are values from 1 to 4294967295 |
| 3277163 | The system cannot automatically specify an "index" for this rule, because a rule with "index" 4294967295 exists. Either specify an unused "index", or update the existing rules so that "index" 4294967295 is not used. |

| Error Code | Description |
| --- | --- |
| 3277149 | The "Anon" field cannot be an empty string |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Export policy ID |
| name | string | Export policy name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_rule

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| policy | policy | |
| protocols | array[string] | |

| Name | Type | Description |
|---|---|---|
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an export policy rule

DELETE `/protocols/nfs/export-policies/{policy.id}/rules/{index}`

**Introduced In:** 9.6

Deletes an export policy rule.

#### Related ONTAP commands

- `vserver export-policy rule delete`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703945 | Ruleset is in use by a volume. It cannot be deleted until all volumes that refer to it are first deleted |
| 1703946 | Cannot determine if the ruleset is in use by a volume. It cannot be deleted until all volumes that refer to it are first deleted |
| 1703954 | Export policy does not exist |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve an export policy rule

GET `/protocols/nfs/export-policies/{policy.id}/rules/{index}`

**Introduced In:** 9.6

Retrieves an export policy rule

#### Related ONTAP commands

- `vserver export-policy rule show`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |

| Name | Type | Description |
| --- | --- | --- |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| policy | policy | |
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anonymous_user": "string",
  "chown_mode": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ],
  "ntfs_unix_security": "string",
  "policy": {
    "name": "string"
  },
  "protocols": [
    "string"
  ],
  "ro_rule": [
    "string"
  ],
  "rw_rule": [
    "string"
  ],
  "superuser": [
    "string"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262197 | The value provided is invalid for the field |
| 6691623 | User is not authorized |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

export_clients

| Name | Type | Description |
|---|---|---|
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

policy

| Name | Type | Description |
|---|---|---|
| id | integer | Export policy ID |
| name | string | Export policy name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the properties of an export policy rule

PATCH `/protocols/nfs/export-policies/{policy.id}/rules/{index}`

**Introduced In:** 9.6

Updates the properties of an export policy rule to change an export policy rule's index or fields.

#### Related ONTAP commands

- `vserver export-policy rule modify`
- `vserver export-policy rule setindex`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |
| new_index | integer | query | False | New Export Rule Index |

#### Request Body

| Name | Type | Description |
|---|---|---|
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |

| Name | Type | Description |
|---|---|---|
| protocols | array[string] | |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "anonymous_user": "string",
  "chown_mode": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ],
  "ntfs_unix_security": "string",
  "protocols": [
    "string"
  ],
  "ro_rule": [
    "string"
  ],
  "rw_rule": [
    "string"
  ],
  "superuser": [
    "string"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262196 | Field 'svm.name' is not supported in the body of PATCH request |
| 262197 | The value provided is invalid for the field |
| 262203 | Field 'svm.uuid' is not supported in the body of PATCH request |
| 1703954 | Export policy does not exist |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704050 | Invalid clientmatch: clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 3277000 | Upgrade all nodes to ONTAP 9.0.0 or above to use krb5p as a security flavor in export-policy rules |

| Error Code | Description |
| --- | --- |
| 3277083 | User ID is not valid. Enter a value for User ID from 0 to 4294967295 |
| 3277149 | The "Anon" field cannot be an empty string |
| 6691623 | User is not authorized |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

export_clients

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |

policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | Export policy ID |
| name | string | Export policy name |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_rule

| Name | Type | Description |
| --- | --- | --- |
| allow_device_creation | boolean | Specifies whether or not device creation is allowed. |
| allow_suid | boolean | Specifies whether or not SetUID bits in SETATTR Op is to be honored. |
| anonymous_user | string | User ID To Which Anonymous Users Are Mapped. |
| chown_mode | string | Specifies who is authorized to change the ownership mode of a file. |
| clients | array[export_clients] | Array of client matches |
| index | integer | Index of the rule within the export policy. |
| ntfs_unix_security | string | NTFS export UNIX security options. |
| protocols | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| ro_rule | array[string] | Authentication flavors that the read-only access rule governs |
| rw_rule | array[string] | Authentication flavors that the read/write access rule governs |
| superuser | array[string] | Authentication flavors that the superuser security type governs |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve export policy rule clients

GET `/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients`

**Introduced In:** 9.6

Retrieves export policy rule clients.

#### Learn more

- DOC /protocols/nfs/export-policies

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of export rule client records |
| records | array[export_client] | |

**Example response**

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "index": 0,
      "match": "0.0.0.0/0",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 6691623 | User is not authorized |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

policy

| Name | Type | Description |
|---|---|---|
| id | integer | Export policy ID |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_client

| Name | Type | Description |
|---|---|---|
| index | integer | Index of the rule within the export policy. |

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |
| policy | policy | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an export policy rule client

POST `/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients`

**Introduced In:** 9.6

Creates an export policy rule client

#### Required properties

- `policy.id` - Existing export policy that contains export policy rules for the client being added.
- `index` - Existing export policy rule for which to create an export client.
- `match` - Base name for the export policy client.

#### Related ONTAP commands

- `vserver export-policy rule add-clientmatches`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| index | integer | Index of the rule within the export policy. |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |
| policy | policy | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "index": 0,
  "match": "0.0.0.0/0",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of export rule client records |
| records | array[export_client] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "index": 0,
      "match": "0.0.0.0/0",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 262196 | Field 'svm.name' is not supported in the body of POST request |
| 1703954 | Export policy does not exist |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704050 | Invalid clientmatch: the clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 6691623 | User is not authorized |

#### Definitions

**See Definitions**

policy

| Name | Type | Description |
|---|---|---|
| id | integer | Export policy ID |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

export_client

| Name | Type | Description |
|---|---|---|
| index | integer | Index of the rule within the export policy. |

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |
| policy | policy | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

export_client

| Name | Type | Description |
| --- | --- | --- |
| index | integer | Index of the rule within the export policy. |

| Name | Type | Description |
| --- | --- | --- |
| match | string | Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:<br>• As a hostname; for instance, host1<br>• As an IPv4 address; for instance, 10.1.12.24<br>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng<br>• As a domain name preceded by the . character; for instance, .example.com |
| policy | policy | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an export policy client

DELETE `/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients/{match}`

**Introduced In:** 9.6

Deletes an export policy client

#### Related ONTAP commands

- `vserver export-policy rule remove-clientmatches`

#### Learn more

- DOC /protocols/nfs/export-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.id | integer | path | True | Export Policy ID |
| index | integer | path | True | Export Rule Index |
| match | string | path | True | Export Client Match |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Response

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1703954 | Export policy does not exist |
| 1704036 | Invalid clientmatch: missing domain name |
| 1704037 | Invalid clientmatch: missing network name |
| 1704038 | Invalid clientmatch: missing netgroup name |
| 1704039 | Invalid clientmatch |
| 1704040 | Invalid clientmatch: address bytes masked out by netmask are non-zero |
| 1704041 | Invalid clientmatch: address bytes masked to zero by netmask |
| 1704042 | Invalid clientmatch: too many bits in netmask |
| 1704043 | Invalid clientmatch: invalid netmask |
| 1704044 | Invalid clientmatch: invalid characters in host name |
| 1704045 | Invalid clientmatch: invalid characters in domain name |
| 1704050 | Invalid clientmatch: the clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported |
| 1704054 | Invalid clientmatch: invalid characters in netgroup name. Valid characters for a netgroup name are 0-9, A-Z, a-z, ".", "_" and "-" |
| 1704064 | Clientmatch host name too long |
| 1704065 | Clientmatch domain name too long |
| 6691623 | User is not authorized |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update Kerberos interfaces

## View and update Kerberos interfaces

### Examples

### Retrieving the Kerberos interface configuration details

```
# The API:
GET /api/protocols/nfs/kerberos/interfaces

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/kerberos/interfaces"
```

### Updating the Kerberos interface configuration

```
# The API:
PATCH /api/protocols/nfs/kerberos/interfaces/{interface.uuid}

# The call:
curl -d "@test_patch_kerb_interface.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/kerberos/interfaces/e62936de-7342-11e8-9eb4-
0050568be2b7"
test_patch_kerb_interface.txt(body):
{
 "enabled" : "true",
 "spn": "nfs/datalif1-vsim3-d1.sim.netapp.com@NFS-NSR-W01.RTP.NETAPP.COM",
 "user" :"administrator",
 "password" :"Hello123!"
}
```

## Retrieve Kerberos interfaces

GET `/protocols/nfs/kerberos/interfaces`

**Introduced In:** 9.6

Retrieves Kerberos interfaces.

### Related ONTAP commands

- `vserver nfs kerberos interface show`

### Learn more

- DOC /protocols/nfs/kerberos/interfaces

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| spn | string | query | False | Filter by spn |
| encryption_types | string | query | False | Filter by encryption_types |
| enabled | boolean | query | False | Filter by enabled |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| machine_account | string | query | False | Filter by machine_account<br>• Introduced in: 9.12 |
| interface.ip.address | string | query | False | Filter by interface.ip.address |
| interface.name | string | query | False | Filter by interface.name |
| interface.uuid | string | query | False | Filter by interface.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[kerberos_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "encryption_types": [
        "string"
      ],
      "force": true,
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "machine_account": "string",
      "spn": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
```

```
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

kerberos_interface

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| enabled | boolean | Specifies if Kerberos is enabled. |
| encryption_types | array[string] | |
| interface | interface | Network interface |
| machine_account | string | Specifies the machine account to create in Active Directory. |
| spn | string | Service principal name. Valid in PATCH. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a Kerberos interface

GET `/protocols/nfs/kerberos/interfaces/{interface.uuid}`

**Introduced In:** 9.6

Retrieves a Kerberos interface.

#### Related ONTAP commands

- `vserver nfs kerberos interface show`

#### Learn more

- DOC /protocols/nfs/kerberos/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.uuid | string | path | True | Network interface UUID<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Specifies if Kerberos is enabled. |
| encryption_types | array[string] | |
| interface | interface | Network interface |

| Name | Type | Description |
| --- | --- | --- |
| machine_account | string | Specifies the machine account to create in Active Directory. |
| spn | string | Service principal name. Valid in PATCH. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "encryption_types": [
    "string"
  ],
  "force": true,
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "machine_account": "string",
  "spn": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Network interface

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update Kerberos interface properties

PATCH `/protocols/nfs/kerberos/interfaces/{interface.uuid}`

**Introduced In:** 9.6

Updates the properties of a Kerberos interface.

#### Optional property

- `force` - Specifies whether the server should ignore any error encountered while deleting the corresponding machine account on the KDC and also disables Kerberos on the LIF. This is applicable only when disabling Kerberos.

#### Related ONTAP commands

- `vserver nfs kerberos interface modify`
- `vserver nfs kerberos interface enable`
- `vserver nfs kerberos interface disable`

#### Learn more

- DOC /protocols/nfs/kerberos/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.uuid | string | path | True | Network interface UUID<br>• Introduced in: 9.7 |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies if Kerberos is enabled. |
| encryption_types | array[string] | |
| force | boolean | Specifies whether the server should ignore any error encountered while deleting the corresponding machine account on the KDC and also disables Kerberos on the LIF. |
| interface | interface | Network interface |
| keytab_uri | string | Load keytab from URI |
| machine_account | string | Specifies the machine account to create in Active Directory. |
| organizational_unit | string | Organizational unit |
| password | string | Account creation password |
| spn | string | Service principal name. Valid in PATCH. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| user | string | Account creation user name |

#### Example request

```
{
  "encryption_types": [
    "string"
  ],
  "interface": {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "keytab_uri": "string",
  "machine_account": "string",
  "organizational_unit": "string",
  "password": "string",
  "spn": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user": "string"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
|---|---|
| 4 | The entry doesn't exist |
| 3276801 | Failed to bind service principal name on LIF. |
| 3276809 | Failed to disable NFS Kerberos on LIF. |
| 3276832 | Failed to insert Kerberos attributes to database. |
| 3276842 | Internal error. Failed to import Kerberos keytab file into the management databases. Contact technical support for assistance. |
| 3276861 | Kerberos is already enabled/disabled on this LIF. |

| Error codes | Description |
| --- | --- |
| 3276862 | Kerberos service principal name is required. |
| 3276889 | Failed to enable NFS Kerberos on LIF. |
| 3276937 | Failed to lookup the Vserver for the virtual interface. |
| 3276941 | Kerberos is a required field. |
| 3276942 | Service principal name is invalid. It must of the format:"nfs/<LIF-FQDN>@REALM"</LIF-FQDN> |
| 3276943 | Lookup of Kerberos realm failed. The entry doesn’t exist. |
| 3276944 | Internal error. Reason: Failed to initialize the Kerberos context |
| 3276945 | Internal error. Reason: Failed to parse the service principal name |
| 3276951 | Warning: Skipping unsupported encryption type for service principal name |
| 3276952 | "organizational_unit" option cannot be used for "Other" vendor. |
| 3276965 | Account sharing across Vservers is not allowed. Use a different service principal name unique within the first 15 characters. |
| 3277019 | Cannot specify -force when enabling Kerberos. |
| 3277020 | Modifying the NFS Kerberos configuration for a LIF that is not configured for NFS is not supported. |
| 3277043 | Keytab import failed due to missing keys. Keys for encryption types are required for Vserver but found no matching keys for service principal name. Generate the keytab file with all required keys and try again. |
| 3277133 | The *machine_account* property can only be specified when the KDC vendor is Microsoft. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip

IP information

interface

Network interface

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

kerberos_interface

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies if Kerberos is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| encryption_types | array[string] | |
| force | boolean | Specifies whether the server should ignore any error encountered while deleting the corresponding machine account on the KDC and also disables Kerberos on the LIF. |
| interface | interface | Network interface |
| keytab_uri | string | Load keytab from URI |
| machine_account | string | Specifies the machine account to create in Active Directory. |
| organizational_unit | string | Organizational unit |
| password | string | Account creation password |
| spn | string | Service principal name. Valid in PATCH. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| user | string | Account creation user name |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage Kerberos realms

### Manage Kerberos realms

#### Examples

### Retrieving the Kerberos realm details

```
# The API:
GET /api/protocols/nfs/kerberos/realms

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/kerberos/realms"
```

### Creating the Kerberos realm for an SVM

```
# The API:
POST /api/protocols/nfs/kerberos/realms

# The call:
curl -d "@test_post_kerb_realm.txt" -X POST "https://<mgmt-
ip>/api/protocols/nfs/kerberos/realms"
test_post_kerb_realm.txt(body):
{
"svm.uuid": "05c90dc2-7343-11e8-9eb4-0050568be2b7",
"name": "NFS-NSR-W02.RTP.NETAPP.COM",
  "kdc": {
    "vendor": "microsoft",
    "ip": "10.225.185.112",
    "port": 88
  },
  "comment": "realm",
  "ad_server": {
    "name": "nfs-nsr-w02.rtp.netapp.com",
    "address": "10.225.185.112"
  }
}
```

#### Updating the Kerberos realm for an SVM

```
# The API:
PATCH /api/protocols/nfs/kerberos/realms/{svm.uuid}/{name}

# The call:
curl -d "@test_patch_kerb_realm.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/kerberos/realms/05c90dc2-7343-11e8-9eb4-
0050568be2b7/NFS-NSR-W02.RTP.NETAPP.COM"
test_patch_kerb_realm.txt(body):
{
"kdc": {
  "vendor": "Microsoft",
  "ip": "100.225.185.112",
  "port": 88
},
"comment": "realm modify",
"ad_server": {
  "name": "nfs.netapp.com",
  "address": "192.2.18.112"
}
}
```

#### Deleting the Kerberos realm for an SVM

```
# The API:
DELETE /api/protocols/nfs/kerberos/realms/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/protocols/nfs/kerberos/realms/05c90dc2-7343-11e8-9eb4-
0050568be2b7/NFS-NSR-W02.RTP.NETAPP.COM"
```

---

### Retrieve Kerberos realms

GET `/protocols/nfs/kerberos/realms`

**Introduced In:** 9.6

Retrieves Kerberos realms.

#### Related ONTAP commands

- `vserver nfs kerberos realm show`

**Learn more**

- DOC /protocols/nfs/kerberos/realms

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| comment | string | query | False | Filter by comment |
| password_server.address | string | query | False | Filter by password_server.address<br>• Introduced in: 9.13 |
| password_server.port | integer | query | False | Filter by password_server.port<br>• Introduced in: 9.13 |
| kdc.vendor | string | query | False | Filter by kdc.vendor |
| kdc.ip | string | query | False | Filter by kdc.ip |
| kdc.port | integer | query | False | Filter by kdc.port<br>• Max value: 65535<br>• Min value: 1 |
| encryption_types | string | query | False | Filter by encryption_types |
| clock_skew | integer | query | False | Filter by clock_skew<br>• Introduced in: 9.13 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| admin_server.address | string | query | False | Filter by admin_server.address<br><br>• Introduced in: 9.13 |
| admin_server.port | integer | query | False | Filter by admin_server.port<br><br>• Introduced in: 9.13 |
| ad_server.name | string | query | False | Filter by ad_server.name |
| ad_server.address | string | query | False | Filter by ad_server.address |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[kerberos_realm] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_server": {
        "address": "1.2.3.4",
        "name": "string"
      },
      "admin_server": {
        "address": "1.2.3.4"
      },
      "comment": "string",
      "encryption_types": [
        "string"
      ],
      "kdc": {
        "ip": "1.2.3.4",
        "port": 88,
        "vendor": "string"
      },
      "name": "string",
      "password_server": {
        "address": "1.2.3.4"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
```

```
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ad_server

| Name | Type | Description |
|---|---|---|
| address | string | Active Directory server IP address |
| name | string | Active Directory server name |

admin_server

| Name | Type | Description |
|---|---|---|
| address | string | Admin server IP address. |
| port | integer | Specifies the port number of admin server. |

kdc

| Name | Type | Description |
|---|---|---|
| ip | string | KDC IP address |
| port | integer | KDC port |

| Name | Type | Description |
| --- | --- | --- |
| vendor | string | Key Distribution Center (KDC) vendor. Following values are supported:<br>• microsoft - Microsoft Active Directory KDC<br>• other - MIT Kerberos KDC or other KDC |

password_server

| Name | Type | Description |
| --- | --- | --- |
| address | string | Password server IP address. |
| port | integer | Specifies the port number of password server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

kerberos_realm

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |

| Name | Type | Description |
|---|---|---|
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a Kerberos realm**

POST `/protocols/nfs/kerberos/realms`

**Introduced In:** 9.6

Creates a Kerberos realm.

**Required properties**

- `svm.uuid` or `svm.name` - Existing SVM on which to create the Kerberos realm.
- `name` - Base name for the Kerberos realm.
- `kdc.vendor` - Vendor of the Key Distribution Center (KDC) server for this Kerberos realm. If the configuration uses a Microsoft Active Directory domain for authentication, this field must be `microsoft`.

- `kdc.ip` - IP address of the KDC server for this Kerberos realm.

#### Recommended optional properties

- `ad_server.name` - Host name of the Active Directory Domain Controller (DC). This is a mandatory parameter if the kdc-vendor is `microsoft`.
- `ad_server.address` - IP address of the Active Directory Domain Controller (DC). This is a mandatory parameter if the kdc-vendor is `microsoft`.

#### Default property values

If not specified in POST, the following default property value is assigned:

- `kdc.port` - *88*
- `admin_server.port` - *749*
- `password_server.port` - *464*
- `clock_skew` - *5*

#### Related ONTAP commands

- `vserver nfs kerberos realm create`

#### Learn more

- DOC /protocols/nfs/kerberos/realms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |

| Name | Type | Description |
|---|---|---|
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "ad_server": {
    "address": "1.2.3.4",
    "name": "string"
  },
  "admin_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": [
    "string"
  ],
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "string"
  },
  "name": "string",
  "password_server": {
    "address": "1.2.3.4"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
|---|---|
| 2949121 | Active Directory server name required. |
| 2949122 | Active Directory server address required |
| 2949123 | Failed to create Kerberos realm. |
| 2949124 | Failed to create hosts file entry. |
| 3276949 | Kerberos realm creation failed. Reason: The parameters "ad_server.name" and "ad_server.address" are only valid when "kdc.vendor" is Microsoft |
| 3276976 | "realm" is a required input |
| 3276998 | Only the data Vservers can own NFS Kerberos realms. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ad_server

| Name | Type | Description |
|---|---|---|
| address | string | Active Directory server IP address |
| name | string | Active Directory server name |

admin_server

| Name | Type | Description |
|---|---|---|
| address | string | Admin server IP address. |
| port | integer | Specifies the port number of admin server. |

kdc

| Name | Type | Description |
|---|---|---|
| ip | string | KDC IP address |
| port | integer | KDC port |
| vendor | string | Key Distribution Center (KDC) vendor. Following values are supported:<br>• microsoft - Microsoft Active Directory KDC<br>• other - MIT Kerberos KDC or other KDC |

password_server

| Name | Type | Description |
| --- | --- | --- |
| address | string | Password server IP address. |
| port | integer | Specifies the port number of password server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

kerberos_realm

| Name | Type | Description |
| --- | --- | --- |
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a Kerberos realm

DELETE `/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`

**Introduced In:** 9.6

Deletes a Kerberos realm.

- `vserver nfs kerberos realm delete`

#### Learn more

- DOC /protocols/nfs/kerberos/realms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| name | string | path | True | Kerberos realm |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
|---|---|
| 1966125 | Failed to remove hosts entry. |
| 1966126 | Failed to lookup hosts entry. |
| 2949141 | Failed to lookup Kerberos realm. |
| 2949142 | Failed to remove Kerberos realm. |
| 3276942 | Service principal name is invalid. It must of the format:"nfs/<LIF-FQDN>@REALM\\\"</LIF-FQDN> |
| 3276976 | "realm" is a required input |
| 3276998 | Only the data Vservers can own NFS Kerberos realms. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a Kerberos realm

GET `/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`

**Introduced In:** 9.6

Retrieves a Kerberos realm.

- `vserver nfs kerberos realm show`

#### Learn more

- DOC /protocols/nfs/kerberos/realms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| name | string | path | True | Kerberos realm |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_server": {
    "address": "1.2.3.4",
    "name": "string"
  },
  "admin_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": [
    "string"
  ],
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "string"
  },
  "name": "string",
  "password_server": {
    "address": "1.2.3.4"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ad_server

| Name | Type | Description |
|---|---|---|
| address | string | Active Directory server IP address |
| name | string | Active Directory server name |

admin_server

| Name | Type | Description |
|---|---|---|
| address | string | Admin server IP address. |
| port | integer | Specifies the port number of admin server. |

kdc

| Name | Type | Description |
|---|---|---|
| ip | string | KDC IP address |
| port | integer | KDC port |
| vendor | string | Key Distribution Center (KDC) vendor. Following values are supported:<br>• microsoft - Microsoft Active Directory KDC<br>• other - MIT Kerberos KDC or other KDC |

password_server

| Name | Type | Description |
| --- | --- | --- |
| address | string | Password server IP address. |
| port | integer | Specifies the port number of password server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update Kerberos realm properties

PATCH `/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`

**Introduced In:** 9.6

Updates the properties of a Kerberos realm.

- `vserver nfs kerberos realm modify`

#### Learn more

- DOC /protocols/nfs/kerberos/realms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| name | string | path | True | Kerberos realm |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "ad_server": {
    "address": "1.2.3.4",
    "name": "string"
  },
  "admin_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": [
    "string"
  ],
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "string"
  },
  "name": "string",
  "password_server": {
    "address": "1.2.3.4"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
|---|---|
| 1966125 | Failed to remove hosts entry. |
| 1966126 | Failed to lookup hosts entry. |

| Error codes | Description |
| --- | --- |
| 1966131 | Failed to create hosts entry. |
| 1966132 | Failed to modify hosts entry. |
| 2949121 | Active Directory server name required. |
| 2949122 | Active Directory server address required |
| 2949123 | Failed to create Kerberos realm. |
| 2949124 | Failed to create hosts file entry. |
| 2949141 | Failed to lookup Kerberos realm. |
| 2949148 | Failed to modify Kerberos realm. |
| 3276976 | "realm" is a required input |
| 3276998 | Only the data Vservers can own NFS Kerberos realms. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ad_server

| Name | Type | Description |
|---|---|---|
| address | string | Active Directory server IP address |
| name | string | Active Directory server name |

admin_server

| Name | Type | Description |
|---|---|---|
| address | string | Admin server IP address. |
| port | integer | Specifies the port number of admin server. |

kdc

| Name | Type | Description |
|---|---|---|
| ip | string | KDC IP address |
| port | integer | KDC port |
| vendor | string | Key Distribution Center (KDC) vendor. Following values are supported:<br>• microsoft - Microsoft Active Directory KDC<br>• other - MIT Kerberos KDC or other KDC |

password_server

| Name | Type | Description |
| --- | --- | --- |
| address | string | Password server IP address. |
| port | integer | Specifies the port number of password server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

kerberos_realm

| Name | Type | Description |
| --- | --- | --- |
| ad_server | ad_server | |
| admin_server | admin_server | |
| clock_skew | integer | Specifies the allowed time of clock-skew between the server and clients, in minutes. |
| comment | string | Comment |
| encryption_types | array[string] | |
| kdc | kdc | |
| name | string | Kerberos realm |
| password_server | password_server | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NFS services

### Manage NFS services

### Retrieving an NFS configuration

```
# The API:
GET /api/protocols/nfs/services

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/services"
```

### Retrieving the mount permissions for a specified volume for a given IP address.

```
# The API:
GET /api/protocols/nfs/services

# The call:
curl -X GET curl -X GET "https://<mgmt-
ip>/api/protocols/nfs/services/{svm.uuid}?protocol_access_rules.volume=tes
tvol12&protocol_access_rules.client_ip=1.2.3.4&protocol_access_rules.auth_
type=sys&fields=protocol_access_rules&return_timeout=15&return_records=tru
e"

#
Returns the protocol_access_rules structure and provides access
permissions for each protocol.
"protocol_access_rules": {
  "nfs3_access_type": "read_write",
  "nfs4_access_type": "read",
  "cifs_access_type": "read_write"
}
```

**Creating an NFS configuration for an SVM**

```
# The API:
POST /api/protocols/nfs/services

# The call:
curl  -d "@test_nfs_post.txt" -X POST "https://<mgmt-
ip>/api/protocols/nfs/services"
test_nfs_post.txt(body):
{
"svm": {
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"protocol": {
  "v4_id_domain": "nfs-nsr-w01.rtp.netapp.com"
},
"vstorage_enabled": "true"
}
```

**Updating an  NFS configuration for an SVM**

```
# The API:
PATCH /api/protocols/nfs/services/{svm.uuid}

# The call:
curl -d "@test_nfs_patch.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/services/4a415601-548c-11e8-a21d-0050568bcbc9"
test_nfs_patch.txt(body):
{
"protocol": {
  "v4_id_domain": "nfs-nsr-w01.rtp.netapp.com"
},
"vstorage_enabled": "false"
}
```

### Deleting an NFS configuration for an SVM

```
# The API:
DELETE /api/protocols/nfs/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/services/4a415601-
548c-11e8-a21d-0050568bcbc9"
```

#### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

### Retrieve the NFS configuration for SVMs

GET `/protocols/nfs/services`

**Introduced In:** 9.6

Retrieves the NFS configuration of SVMs.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Advanced properties

- security.rpcsec_context_idle
- security.ntfs_unix_security
- security.chown_mode
- security.nt_acl_display_permission
- protocol.v3_features.ejukebox_enabled
- protocol.v3_features.connection_drop
- protocol.v3_features.fsid_change
- protocol.v3_features.mount_daemon_port
- protocol.v3_features.network_lock_manager_port
- protocol.v3_features.network_status_monitor_port
- protocol.v3_features.rquota_daemon_port
- protocol.v3_features.hide_snapshot_enabled
- protocol.v41_features.implementation_domain
- protocol.v41_features.implementation_name
- protocol.v40_features.acl_max_aces
- windows.map_unknown_uid_to_default_user
- exports.netgroup_trust_any_nsswitch_no_match
- credential_cache.negative_ttl
- transport.tcp_max_transfer_size
- root.*
- protocol.v41_features.trunking_enabled
- protocol.v42_features.seclabel_enabled
- protocol.v42_features.sparsefile_ops_enabled
- protocol.v42_features.xattrs_enabled
- protocol.v40_features.referrals_enabled
- protocol.v41_features.referrals_enabled
- protocol.v4_lease_seconds
- protocol.v4_session_slots
- protocol.v4_session_slot_reply_cache_size
- protocol.v4_fsid_change

#### Diagnostic properties

- credential_cache.transient_error_ttl

- `access_cache_config.ttl_failure`
- `credential_cache.harvest_timeout`

#### Related ONTAP commands

- `vserver nfs show`
- `vserver nfs status`
- `vserver export-policy check-access`

#### Learn more

- DOC /protocols/nfs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol_access_rules.volume | string | query | False | Volume on which access needs to be checked.<br>• Introduced in: 9.10 |
| protocol_access_rules.client_ip | string | query | False | IP address for the client for which access needs to be checked.<br>• Introduced in: 9.10 |
| protocol_access_rules.auth_type | string | query | False | Authentication method used to check the client's access to the volume.<br>• Introduced in: 9.10<br>• Default value: 1 |
| root.ignore_nt_acl | boolean | query | False | Filter by root.ignore_nt_acl<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| root.skip_write_permission_check | boolean | query | False | Filter by root.skip_write_permission_check<br>• Introduced in: 9.11 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| file_session_io_grouping_duration | integer | query | False | Filter by file_session_io_grouping_duration<br>• Introduced in: 9.11<br>• Max value: 3600<br>• Min value: 60 |
| extended_groups_limit | integer | query | False | Filter by extended_groups_limit<br>• Introduced in: 9.8<br>• Max value: 1024<br>• Min value: 32 |
| protocol_access_rules.nfs3_access_type | string | query | False | Filter by protocol_access_rules.nfs3_access_type<br>• Introduced in: 9.10 |
| protocol_access_rules.cifs_access_type | string | query | False | Filter by protocol_access_rules.cifs_access_type<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol_access_rules.nfs4_access_type | string | query | False | Filter by protocol_access_rules.nfs4_access_type<br>• Introduced in: 9.10 |
| transport.tcp_enabled | boolean | query | False | Filter by transport.tcp_enabled |
| transport.udp_enabled | boolean | query | False | Filter by transport.udp_enabled |
| transport.rdma_enabled | boolean | query | False | Filter by transport.rdma_enabled<br>• Introduced in: 9.12 |
| transport.tcp_max_transfer_size | integer | query | False | Filter by transport.tcp_max_transfer_size<br>• Introduced in: 9.11<br>• Max value: 1048576<br>• Min value: 8192 |
| security.permitted_encryption_types | string | query | False | Filter by security.permitted_encryption_types<br>• Introduced in: 9.11 |
| security.ntfs_unix_security | string | query | False | Filter by security.ntfs_unix_security<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security.rpcsec_context_idle | integer | query | False | Filter by security.rpcsec_context_idle<br><br>• Introduced in: 9.11 |
| security.chown_mode | string | query | False | Filter by security.chown_mode<br><br>• Introduced in: 9.11 |
| security.nt_acl_display_permission | boolean | query | False | Filter by security.nt_acl_display_permission<br><br>• Introduced in: 9.11 |
| qtree.validate_export | boolean | query | False | Filter by qtree.validate_export<br><br>• Introduced in: 9.10 |
| qtree.export_enabled | boolean | query | False | Filter by qtree.export_enabled<br><br>• Introduced in: 9.10 |
| auth_sys_extended_groups_enabled | boolean | query | False | Filter by auth_sys_extended_groups_enabled<br><br>• Introduced in: 9.8 |
| showmount_enabled | boolean | query | False | Filter by showmount_enabled<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rquota_enabled | boolean | query | False | Filter by rquota_enabled<br><br>• Introduced in: 9.8 |
| statistics.v3.latency_raw.other | integer | query | False | Filter by statistics.v3.latency_raw.other<br><br>• Introduced in: 9.7 |
| statistics.v3.latency_raw.total | integer | query | False | Filter by statistics.v3.latency_raw.total<br><br>• Introduced in: 9.7 |
| statistics.v3.latency_raw.write | integer | query | False | Filter by statistics.v3.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.v3.latency_raw.read | integer | query | False | Filter by statistics.v3.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.v3.iops_raw.other | integer | query | False | Filter by statistics.v3.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.v3.iops_raw.total | integer | query | False | Filter by statistics.v3.iops_raw.total<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.v3.iops_raw.write | integer | query | False | Filter by statistics.v3.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.v3.iops_raw.read | integer | query | False | Filter by statistics.v3.iops_raw.read<br><br>• Introduced in: 9.7 |
| statistics.v3.status | string | query | False | Filter by statistics.v3.status<br><br>• Introduced in: 9.7 |
| statistics.v3.throughput_raw.write | integer | query | False | Filter by statistics.v3.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.v3.throughput_raw.read | integer | query | False | Filter by statistics.v3.throughput_raw.read<br><br>• Introduced in: 9.7 |
| statistics.v3.throughput_raw.total | integer | query | False | Filter by statistics.v3.throughput_raw.total<br><br>• Introduced in: 9.7 |
| statistics.v3.timestamp | string | query | False | Filter by statistics.v3.timestamp<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.v41.latency_raw.other | integer | query | False | Filter by statistics.v41.latency_raw.other<br><br>• Introduced in: 9.8 |
| statistics.v41.latency_raw.total | integer | query | False | Filter by statistics.v41.latency_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v41.latency_raw.write | integer | query | False | Filter by statistics.v41.latency_raw.write<br><br>• Introduced in: 9.8 |
| statistics.v41.latency_raw.read | integer | query | False | Filter by statistics.v41.latency_raw.read<br><br>• Introduced in: 9.8 |
| statistics.v41.iops_raw.other | integer | query | False | Filter by statistics.v41.iops_raw.other<br><br>• Introduced in: 9.8 |
| statistics.v41.iops_raw.total | integer | query | False | Filter by statistics.v41.iops_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v41.iops_raw.write | integer | query | False | Filter by statistics.v41.iops_raw.write<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.v41.iops_raw.read | integer | query | False | Filter by statistics.v41.iops_raw.read<br><br>• Introduced in: 9.8 |
| statistics.v41.status | string | query | False | Filter by statistics.v41.status<br><br>• Introduced in: 9.8 |
| statistics.v41.throughput_raw.write | integer | query | False | Filter by statistics.v41.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.v41.throughput_raw.read | integer | query | False | Filter by statistics.v41.throughput_raw.read<br><br>• Introduced in: 9.8 |
| statistics.v41.throughput_raw.total | integer | query | False | Filter by statistics.v41.throughput_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v41.timestamp | string | query | False | Filter by statistics.v41.timestamp<br><br>• Introduced in: 9.8 |
| statistics.v4.latency_raw.other | integer | query | False | Filter by statistics.v4.latency_raw.other<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.v4.latency_raw.total | integer | query | False | Filter by statistics.v4.latency_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v4.latency_raw.write | integer | query | False | Filter by statistics.v4.latency_raw.write<br><br>• Introduced in: 9.8 |
| statistics.v4.latency_raw.read | integer | query | False | Filter by statistics.v4.latency_raw.read<br><br>• Introduced in: 9.8 |
| statistics.v4.iops_raw.other | integer | query | False | Filter by statistics.v4.iops_raw.other<br><br>• Introduced in: 9.8 |
| statistics.v4.iops_raw.total | integer | query | False | Filter by statistics.v4.iops_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v4.iops_raw.write | integer | query | False | Filter by statistics.v4.iops_raw.write<br><br>• Introduced in: 9.8 |
| statistics.v4.iops_raw.read | integer | query | False | Filter by statistics.v4.iops_raw.read<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.v4.status | string | query | False | Filter by statistics.v4.status<br><br>• Introduced in: 9.8 |
| statistics.v4.throughput_raw.write | integer | query | False | Filter by statistics.v4.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.v4.throughput_raw.read | integer | query | False | Filter by statistics.v4.throughput_raw.read<br><br>• Introduced in: 9.8 |
| statistics.v4.throughput_raw.total | integer | query | False | Filter by statistics.v4.throughput_raw.total<br><br>• Introduced in: 9.8 |
| statistics.v4.timestamp | string | query | False | Filter by statistics.v4.timestamp<br><br>• Introduced in: 9.8 |
| state | string | query | False | Filter by state |
| enabled | boolean | query | False | Filter by enabled |
| exports.name_service_lookup_protocol | string | query | False | Filter by exports.name_service_lookup_protocol<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| exports.netgroup_trust_any_nsswitch_no_match | boolean | query | False | Filter by exports.netgroup_trust_any_nsswitch_no_match<br>• Introduced in: 9.11 |
| metric.v41.iops.other | integer | query | False | Filter by metric.v41.iops.other<br>• Introduced in: 9.8 |
| metric.v41.iops.total | integer | query | False | Filter by metric.v41.iops.total<br>• Introduced in: 9.8 |
| metric.v41.iops.write | integer | query | False | Filter by metric.v41.iops.write<br>• Introduced in: 9.8 |
| metric.v41.iops.read | integer | query | False | Filter by metric.v41.iops.read<br>• Introduced in: 9.8 |
| metric.v41.latency.other | integer | query | False | Filter by metric.v41.latency.other<br>• Introduced in: 9.8 |
| metric.v41.latency.total | integer | query | False | Filter by metric.v41.latency.total<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.v41.latency.write | integer | query | False | Filter by metric.v41.latency.write<br><br>• Introduced in: 9.8 |
| metric.v41.latency.read | integer | query | False | Filter by metric.v41.latency.read<br><br>• Introduced in: 9.8 |
| metric.v41.timestamp | string | query | False | Filter by metric.v41.timestamp<br><br>• Introduced in: 9.8 |
| metric.v41.throughput.write | integer | query | False | Filter by metric.v41.throughput.write<br><br>• Introduced in: 9.8 |
| metric.v41.throughput.read | integer | query | False | Filter by metric.v41.throughput.read<br><br>• Introduced in: 9.8 |
| metric.v41.throughput.total | integer | query | False | Filter by metric.v41.throughput.total<br><br>• Introduced in: 9.8 |
| metric.v41.status | string | query | False | Filter by metric.v41.status<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.v41.duration | string | query | False | Filter by metric.v41.duration<br><br>• Introduced in: 9.8 |
| metric.v3.iops.other | integer | query | False | Filter by metric.v3.iops.other<br><br>• Introduced in: 9.7 |
| metric.v3.iops.total | integer | query | False | Filter by metric.v3.iops.total<br><br>• Introduced in: 9.7 |
| metric.v3.iops.write | integer | query | False | Filter by metric.v3.iops.write<br><br>• Introduced in: 9.7 |
| metric.v3.iops.read | integer | query | False | Filter by metric.v3.iops.read<br><br>• Introduced in: 9.7 |
| metric.v3.latency.other | integer | query | False | Filter by metric.v3.latency.other<br><br>• Introduced in: 9.7 |
| metric.v3.latency.total | integer | query | False | Filter by metric.v3.latency.total<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.v3.latency.write | integer | query | False | Filter by metric.v3.latency.write<br><br>• Introduced in: 9.7 |
| metric.v3.latency.read | integer | query | False | Filter by metric.v3.latency.read<br><br>• Introduced in: 9.7 |
| metric.v3.timestamp | string | query | False | Filter by metric.v3.timestamp<br><br>• Introduced in: 9.7 |
| metric.v3.throughput.write | integer | query | False | Filter by metric.v3.throughput.write<br><br>• Introduced in: 9.7 |
| metric.v3.throughput.read | integer | query | False | Filter by metric.v3.throughput.read<br><br>• Introduced in: 9.7 |
| metric.v3.throughput.total | integer | query | False | Filter by metric.v3.throughput.total<br><br>• Introduced in: 9.7 |
| metric.v3.status | string | query | False | Filter by metric.v3.status<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.v3.duration | string | query | False | Filter by metric.v3.duration<br>• Introduced in: 9.7 |
| metric.v4.iops.other | integer | query | False | Filter by metric.v4.iops.other<br>• Introduced in: 9.8 |
| metric.v4.iops.total | integer | query | False | Filter by metric.v4.iops.total<br>• Introduced in: 9.8 |
| metric.v4.iops.write | integer | query | False | Filter by metric.v4.iops.write<br>• Introduced in: 9.8 |
| metric.v4.iops.read | integer | query | False | Filter by metric.v4.iops.read<br>• Introduced in: 9.8 |
| metric.v4.latency.other | integer | query | False | Filter by metric.v4.latency.other<br>• Introduced in: 9.8 |
| metric.v4.latency.total | integer | query | False | Filter by metric.v4.latency.total<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.v4.latency.write | integer | query | False | Filter by metric.v4.latency.write<br><br>• Introduced in: 9.8 |
| metric.v4.latency.read | integer | query | False | Filter by metric.v4.latency.read<br><br>• Introduced in: 9.8 |
| metric.v4.timestamp | string | query | False | Filter by metric.v4.timestamp<br><br>• Introduced in: 9.8 |
| metric.v4.throughput.write | integer | query | False | Filter by metric.v4.throughput.write<br><br>• Introduced in: 9.8 |
| metric.v4.throughput.read | integer | query | False | Filter by metric.v4.throughput.read<br><br>• Introduced in: 9.8 |
| metric.v4.throughput.total | integer | query | False | Filter by metric.v4.throughput.total<br><br>• Introduced in: 9.8 |
| metric.v4.status | string | query | False | Filter by metric.v4.status<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.v4.duration | string | query | False | Filter by metric.v4.duration<br><br>• Introduced in: 9.8 |
| windows.default_user | string | query | False | Filter by windows.default_user<br><br>• Introduced in: 9.11 |
| windows.v3_ms_dos_client_enabled | boolean | query | False | Filter by windows.v3_ms_dos_client_enabled<br><br>• Introduced in: 9.11 |
| windows.map_unknown_uid_to_default_user | boolean | query | False | Filter by windows.map_unknown_uid_to_default_user<br><br>• Introduced in: 9.11 |
| access_cache_config.ttl_failure | integer | query | False | Filter by access_cache_config.ttl_failure<br><br>• Introduced in: 9.10 |
| access_cache_config.ttl_negative | integer | query | False | Filter by access_cache_config.ttl_negative<br><br>• Introduced in: 9.10 |
| access_cache_config.isDnsTTLEnabled | boolean | query | False | Filter by access_cache_config.isDnsTTLEnabled<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| access_cache_config.harvest_timeout | integer | query | False | Filter by access_cache_config.harvest_timeout<br><br>• Introduced in: 9.10 |
| access_cache_config.ttl_positive | integer | query | False | Filter by access_cache_config.ttl_positive<br><br>• Introduced in: 9.10 |
| protocol.v4_fsid_change | boolean | query | False | Filter by protocol.v4_fsid_change<br><br>• Introduced in: 9.13 |
| protocol.v4_id_domain | string | query | False | Filter by protocol.v4_id_domain |
| protocol.v40_enabled | boolean | query | False | Filter by protocol.v40_enabled |
| protocol.v4_grace_seconds | integer | query | False | Filter by protocol.v4_grace_seconds<br><br>• Introduced in: 9.13 |
| protocol.v3_64bit_identifiers_enabled | boolean | query | False | Filter by protocol.v3_64bit_identifiers_enabled<br><br>• Introduced in: 9.8 |
| protocol.v4_64bit_identifiers_enabled | boolean | query | False | Filter by protocol.v4_64bit_identifiers_enabled<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| protocol.v4_lease_seconds | integer | query | False | Filter by protocol.v4_lease_seconds<br><br>• Introduced in: 9.13 |
| protocol.v4_subnet_filter_enabled | boolean | query | False | Filter by protocol.v4_subnet_filter_enabled<br><br>• Introduced in: 9.18 |
| protocol.v3_enabled | boolean | query | False | Filter by protocol.v3_enabled |
| protocol.v40_features.acl_preserve | boolean | query | False | Filter by protocol.v40_features.acl_preserve<br><br>• Introduced in: 9.11 |
| protocol.v40_features.write_delegation_enabled | boolean | query | False | Filter by protocol.v40_features.write_delegation_enabled |
| protocol.v40_features.read_delegation_enabled | boolean | query | False | Filter by protocol.v40_features.read_delegation_enabled |
| protocol.v40_features.acl_enabled | boolean | query | False | Filter by protocol.v40_features.acl_enabled |
| protocol.v40_features.referrals_enabled | boolean | query | False | Filter by protocol.v40_features.referrals_enabled<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol.v40_features.acl_max_aces | integer | query | False | Filter by protocol.v40_features.acl_max_aces<br>• Introduced in: 9.11<br>• Max value: 1024<br>• Min value: 192 |
| protocol.v4_session_slot_reply_cache_size | integer | query | False | Filter by protocol.v4_session_slot_reply_cache_size<br>• Introduced in: 9.13<br>• Max value: 4096<br>• Min value: 512 |
| protocol.v41_enabled | boolean | query | False | Filter by protocol.v41_enabled |
| protocol.v41_features.pnfs_enabled | boolean | query | False | Filter by protocol.v41_features.pnfs_enabled |
| protocol.v41_features.read_delegation_enabled | boolean | query | False | Filter by protocol.v41_features.read_delegation_enabled |
| protocol.v41_features.acl_enabled | boolean | query | False | Filter by protocol.v41_features.acl_enabled |
| protocol.v41_features.write_delegation_enabled | boolean | query | False | Filter by protocol.v41_features.write_delegation_enabled |
| protocol.v41_features.trunking_enabled | boolean | query | False | Filter by protocol.v41_features.trunking_enabled<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol.v41_features.referrals_enabled | boolean | query | False | Filter by protocol.v41_features.referrals_enabled<br><br>• Introduced in: 9.13 |
| protocol.v41_features.implementation_name | string | query | False | Filter by protocol.v41_features.implementation_name<br><br>• Introduced in: 9.11 |
| protocol.v41_features.implementation_domain | string | query | False | Filter by protocol.v41_features.implementation_domain<br><br>• Introduced in: 9.11 |
| protocol.v3_features.network_status_monitor_port | integer | query | False | Filter by protocol.v3_features.network_status_monitor_port<br><br>• Introduced in: 9.11 |
| protocol.v3_features.fsid_change | boolean | query | False | Filter by protocol.v3_features.fsid_change<br><br>• Introduced in: 9.11 |
| protocol.v3_features.ejukebox_enabled | boolean | query | False | Filter by protocol.v3_features.ejukebox_enabled<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol.v3_features.mount_daemon_port | integer | query | False | Filter by protocol.v3_features.mount_daemon_port<br><br>• Introduced in: 9.11 |
| protocol.v3_features.rquota_daemon_port | integer | query | False | Filter by protocol.v3_features.rquota_daemon_port<br><br>• Introduced in: 9.11 |
| protocol.v3_features.hide_snapshot_enabled | boolean | query | False | Filter by protocol.v3_features.hide_snapshot_enabled<br><br>• Introduced in: 9.13 |
| protocol.v3_features.connection_drop | boolean | query | False | Filter by protocol.v3_features.connection_drop<br><br>• Introduced in: 9.11 |
| protocol.v3_features.mount_root_only | boolean | query | False | Filter by protocol.v3_features.mount_root_only<br><br>• Introduced in: 9.11 |
| protocol.v3_features.network_lock_manager_port | integer | query | False | Filter by protocol.v3_features.network_lock_manager_port<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| protocol.v42_features.xattrs_enabled | boolean | query | False | Filter by protocol.v42_features.xattrs_enabled<br>• Introduced in: 9.12 |
| protocol.v42_features.sparsefile_ops_enabled | boolean | query | False | Filter by protocol.v42_features.sparsefile_ops_enabled<br>• Introduced in: 9.12 |
| protocol.v42_features.seclabel_enabled | boolean | query | False | Filter by protocol.v42_features.seclabel_enabled<br>• Introduced in: 9.12 |
| protocol.v4_session_slots | integer | query | False | Filter by protocol.v4_session_slots<br>• Introduced in: 9.13<br>• Max value: 2000<br>• Min value: 1 |
| credential_cache.harvest_timeout | integer | query | False | Filter by credential_cache.harvest_timeout<br>• Introduced in: 9.18<br>• Max value: 604800000<br>• Min value: 60000 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| credential_cache.transient_error_ttl | integer | query | False | Filter by credential_cache.transient_error_ttl<br>• Introduced in: 9.11<br>• Max value: 300000<br>• Min value: 30000 |
| credential_cache.negative_ttl | integer | query | False | Filter by credential_cache.negative_ttl<br>• Introduced in: 9.11<br>• Max value: 604800000<br>• Min value: 60000 |
| credential_cache.positive_ttl | integer | query | False | Filter by credential_cache.positive_ttl<br>• Introduced in: 9.11<br>• Max value: 604800000<br>• Min value: 60000 |
| file_session_io_grouping_count | integer | query | False | Filter by file_session_io_grouping_count<br>• Introduced in: 9.11<br>• Max value: 120000<br>• Min value: 1000 |
| vstorage_enabled | boolean | query | False | Filter by vstorage_enabled |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of NFS Server Records |
| records | array[nfs_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_cache_config": {
        "harvest_timeout": 3600,
        "ttl_failure": 1,
        "ttl_negative": 30,
        "ttl_positive": 60
      },
      "credential_cache": {
        "harvest_timeout": 86400000,
        "negative_ttl": 3600000,
        "positive_ttl": 3600000,
        "transient_error_ttl": 30000
      },
      "exports": {
        "name_service_lookup_protocol": "string"
      },
      "extended_groups_limit": 32,
      "file_session_io_grouping_count": 5000,
      "file_session_io_grouping_duration": 120,
      "metric": {
        "v3": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "duration": "PT15S",
          "iops": {
```

```
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "v4": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "v41": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
            }
          },
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        }
      },
      "protocol": {
        "v40_features": {
          "acl_max_aces": 500
        },
        "v41_features": {
          "implementation_domain": "string",
          "implementation_name": "string"
        },
        "v4_id_domain": "string"
      },
      "protocol_access_rules": {
        "cifs_access_type": "string",
        "nfs3_access_type": "string",
        "nfs4_access_type": "string"
      },
      "security": {
        "chown_mode": "string",
        "ntfs_unix_security": "string",
        "permitted_encryption_types": [
          "string"
        ]
      },
      "state": "string",
      "statistics": {
```

```
        "v3": {
          "iops_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "v4": {
          "iops_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "v41": {
          "iops_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency_raw": {
            "read": 200,
```

```
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "transport": {
        "tcp_max_transfer_size": 16384
      },
      "windows": {
        "default_user": "string"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_cache_config

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the time after which an entry is deleted from the access cache, if unused. |
| isDnsTTLEnabled | boolean | Specifies whether Dns TTL is enabled. |
| ttl_failure | integer | Specifies the time to live value for entries for which a failure was encountered, in seconds. |
| ttl_negative | integer | Specifies the time to live value of a negative access cache, in seconds. |
| ttl_positive | integer | Specifies the time to live value of a positive access cache, in seconds. |

credential_cache

| Name | Type | Description |
| --- | --- | --- |
| harvest_timeout | integer | Specifies the age of the unused cached entries, in milliseconds, after which they are cleared from the cache. |
| negative_ttl | integer | Specifies the age in milliseconds, of the negative cached credentials after which they are cleared from the cache. |
| positive_ttl | integer | Specifies the age in milliseconds, of the positive cached credentials after which they are cleared from the cache. |
| transient_error_ttl | integer | Specifies the age in milliseconds, of the cached entries during a transient error situation. |

exports

| Name | Type | Description |
| --- | --- | --- |
| name_service_lookup_protocol | string | Specifies the protocol to use for doing name service lookups. |
| netgroup_trust_any_nsswitch_no_match | boolean | Specifies if you can consider a no-match result from any of the netgroup ns-switch sources to be authoritative. If this option is enabled, then a no-match response from any of the netgroup ns-switch sources is deemed conclusive even if other sources could not be searched. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
| --- | --- | --- |
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

v3_features

| Name | Type | Description |
| --- | --- | --- |
| connection_drop | boolean | Specifies whether the dropping of a connection when an NFSv3 request is dropped is enabled. |
| ejukebox_enabled | boolean | Specifies whether NFSv3 EJUKEBOX error is enabled. |
| fsid_change | boolean | Specifies whether the change in FSID as NFSv3 clients traverse filesystems should be shown. |
| hide_snapshot_enabled | boolean | Specifies whether hiding a snapshot directory under a NFSv3 mount point is enabled. |
| mount_daemon_port | integer | Specifies which port the NFS mount daemon (mountd) uses. |
| mount_root_only | boolean | Specifies whether the SVM allows MOUNT protocol calls only from privileged ports (port numbers less than 1024). |

| Name | Type | Description |
| --- | --- | --- |
| network_lock_manager_port | integer | Specifies which port the Network lock manager uses. |
| network_status_monitor_port | integer | Specifies which port the Network status monitor port uses. |
| rquota_daemon_port | integer | Specifies which port the NFS quota daemon port uses. |

v40_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.0 ACLs is enabled. |
| acl_max_aces | integer | Specifies the maximum number of aces in a NFSv4.0 ACL. |
| acl_preserve | boolean | Specifies if the NFSv4 ACL is preserved or dropped when chmod is performed. In unified security style, this parameter also specifies if NTFS file permissions are preserved or dropped when chmod, chgrp, or chown are performed. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.0 Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.0 referrals is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.0 Write Delegation is enabled. |

v41_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.1 or later ACLs is enabled. |
| implementation_domain | string | Specifies the NFSv4.1 or later implementation ID domain. |

| Name | Type | Description |
| --- | --- | --- |
| implementation_name | string | Specifies the NFSv4.1 or later implementation ID name. |
| pnfs_enabled | boolean | Specifies whether NFSv4.1 or later Parallel NFS is enabled. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.1 referrals is enabled. |
| trunking_enabled | boolean | Specifies whether NFSv4.1 or later trunking is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Write Delegation is enabled. |

v42_features

| Name | Type | Description |
| --- | --- | --- |
| seclabel_enabled | boolean | Specifies whether NFSv4.2 or later security label is enabled. |
| sparsefile_ops_enabled | boolean | Specifies whether NFSv4.2 or later sparsefile operation is enabled. |
| xattrs_enabled | boolean | Specifies whether NFSv4.2 or later extended attributes is enabled. |

protocol

| Name | Type | Description |
| --- | --- | --- |
| v3_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled. |
| v3_enabled | boolean | Specifies whether NFSv3 protocol is enabled. |
| v3_features | v3_features | |

| Name | Type | Description |
| --- | --- | --- |
| v40_enabled | boolean | Specifies whether NFSv4.0 protocol is enabled. |
| v40_features | v40_features | |
| v41_enabled | boolean | Specifies whether NFSv4.1 or later protocol is enabled. |
| v41_features | v41_features | |
| v42_features | v42_features | |
| v4_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled. |
| v4_fsid_change | boolean | Specifies whether the change in FSID when NFSv4 clients traverse file systems is displayed. |
| v4_grace_seconds | integer | Specifies the grace period for clients to reclaim file locks after a server failure. |
| v4_id_domain | string | Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol. |
| v4_lease_seconds | integer | Specifies the lease seconds of the NFSv4 clients. If it is inactive for more than the time displayed, all of the file lock states on a node might be lost. |
| v4_session_slot_reply_cache_size | integer | Specifies the number of bytes of the reply that is cached in each NFSv4.x session slot. |
| v4_session_slots | integer | Specifies the number of entries in NFSv4.x session slot table. |
| v4_subnet_filter_enabled | boolean | Specifies whether NFSv4 subnet filtering is enabled. |

protocol_access_rules

| Name | Type | Description |
| --- | --- | --- |
| cifs_access_type | string | Access available for the CIFS protocol. |
| nfs3_access_type | string | Access available for the NFSv3 protocol. |
| nfs4_access_type | string | Access available for the NFSv4 protocol. |

qtree

| Name | Type | Description |
| --- | --- | --- |
| export_enabled | boolean | Specifies whether qtree export is enabled. |
| validate_export | boolean | Specifies whether qtree export validation is enabled. |

root

| Name | Type | Description |
| --- | --- | --- |
| ignore_nt_acl | boolean | Specifies whether Windows ACLs affect root access from NFS. If this option is enabled, root access from NFS ignores the NT ACL set on the file or directory. |
| skip_write_permission_check | boolean | Specifies if permission checks are to be skipped for NFS WRITE calls from root/owner. For copying read-only files to a destination folder which has inheritable ACLs, this option must be enabled. |

security

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies whether file ownership can be changed only by the superuser, or if a non-root user can also change file ownership. If you set this parameter to restricted, file ownership can be changed only by the superuser, even though the on-disk permissions allow a non-root user to change file ownership. If you set this parameter to unrestricted, file ownership can be changed by the superuser and by the non-root user, depending upon the access granted by on-disk permissions. If you set this parameter to use_export_policy, file ownership can be changed in accordance with the relevant export rules. |
| nt_acl_display_permission | boolean | Controls the permissions that are displayed to NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set. When true, the displayed permissions are based on the maximum access granted by the NT ACL to any user. When false, the displayed permissions are based on the minimum access granted by the NT ACL to any user. |

| Name | Type | Description |
| --- | --- | --- |
| ntfs_unix_security | string | Specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to ignore, ONTAP ignores NFSv3 security changes. If you set this parameter to fail, this overrides the UNIX security options set in the relevant export rules. If you set this parameter to use_export_policy, ONTAP processes NFSv3 security changes in accordance with the relevant export rules.<br>• Default value: 1<br>• enum: ["ignore", "fail", "use_export_policy"]<br>• Introduced in: 9.11<br>• x-nullable: true |
| permitted_encryption_types | array[string] | Specifies the permitted encryption types for Kerberos over NFS. |
| rpcsec_context_idle | integer | Specifies, in seconds, the amount of time a RPCSEC_GSS context is permitted to remain unused before it is deleted. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

statistics

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
| --- | --- | --- |
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

transport

| Name | Type | Description |
| --- | --- | --- |
| rdma_enabled | boolean | Specifies whether RDMA transport is enabled for the NFS server. |
| tcp_enabled | boolean | Specifies whether TCP transports are enabled on the server. |

| Name | Type | Description |
| --- | --- | --- |
| tcp_max_transfer_size | integer | Specifies the maximum transfer size in bytes, that the storage system negotiates with the client for TCP transport of data for NFSv3 and NFSv4.x protocols. The range is 8192 to 1048576. |
| udp_enabled | boolean | Specifies whether UDP transports are enabled on the server. |

windows

| Name | Type | Description |
| --- | --- | --- |
| default_user | string | Specifies the default Windows user for the NFS server. |
| map_unknown_uid_to_default_user | boolean | Specifies whether or not the mapping of an unknown UID to the default Windows user is enabled. |
| v3_ms_dos_client_enabled | boolean | Specifies whether NFSv3 MS-DOS client support is enabled. |

nfs_service

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |

| Name | Type | Description |
| --- | --- | --- |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| metric | metric | Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>* online - NFS server is ready to accept client requests.<br>* offline - NFS server is not ready to accept client requests. |
| statistics | statistics | Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |

| Name | Type | Description |
|---|---|---|
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an NFS configuration for an SVM**

POST `/protocols/nfs/services`

**Introduced In:** 9.6

Creates an NFS configuration for an SVM.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM for which to create the NFS configuration.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `state` - online
- `transport.udp_enabled` - *true*

- `transport.tcp_enabled` - *true*
- `transport.rdma_enabled` - *true*
- `protocol.v3_enabled` - *true*
- `protocol.v3_64bit_identifiers_enabled` - *false*
- `protocol.v4_id_domain` - defaultv4iddomain.com
- `protocol.v4_64bit_identifiers_enabled` - *true*
- `protocol.v4_enabled` - *false*
- `protocol.v41_enabled` - *false*
- `protocol.v40_features.acl_enabled` - *false*
- `protocol.v40_features.read_delegation_enabled` - *false*
- `protocol.v40_features.write_delegation_enabled` - *false*
- `protocol.v41_features.acl_enabled` - *false*
- `protocol.v41_features.read_delegation_enabled` - *false*
- `protocol.v41_features.write_delegation_enabled` - *false*
- `protocol.v41_features.pnfs_enabled` - *false*
- `protocol.v4_features.subnet_filter_enabled` - *false*
- `vstorage_enabled` - *false*
- `rquota_enabled` - *false*
- `showmount_enabled` - *true*
- `auth_sys_extended_groups_enabled` - *false*
- `extended_groups_limit` - *32*
- `qtree.export_enabled` - *false*
- `qtree.validate_export` - *true*
- `access_cache_config.ttl_positive` - *60*
- `access_cache_config.ttl_negative` - *30*
- `access_cache_config.ttl_failure` - *1*
- `access_cache_config.harvest_timeout` - *3600*
- `access_cache_config.isDnsTTLEnabled` - *false*
- `file_session_io_grouping_count` - *5000*
- `file_session_io_grouping_duration` - *120*
- `security.nt_acl_display_permission` - *false*
- `exports.netgroup_trust_any_nsswitch_no_match` - *false*
- `exports.name_service_lookup_protocol` - *udp*
- `security.permitted_encryption_types` - [aes-256,aes-128,des3,des]

- security.rpcsec_context_idle - *0*
- security.chown_mode - *use_export_policy*
- security.ntfs_unix_security - *use_export_policy*
- windows.v3_ms_dos_client_enabled - *false*
- windows.default_user - ""
- windows.map_unknown_uid_to_default_user - *true*
- credential_cache.positive_ttl - *3600000*
- credential_cache.negative_ttl - *3600000*
- credential_cache.transient_error_ttl - *30000*
- credential_cache.harvest_timeout - *86400000*
- protocol.v40_features.acl_preserve - *true*
- protocol.v41_features.implementation_domain - 'netapp.com'
- protocol.v40_features.acl_max_aces - *400*
- protocol.v3_features.ejukebox_enabled - *true*
- protocol.v3_features.connection_drop - *true*
- protocol.v3_features.fsid_change - *true*
- protocol.v3_features.mount_daemon_port - *635*
- protocol.v3_features.network_lock_manager_port - *4045*
- protocol.v3_features.network_status_monitor_port - *4046*
- protocol.v3_features.rquota_daemon_port - *4046*
- protocol.v3_features.mount_root_only - *true*
- transport.tcp_max_transfer_size - *65536*
- root.ignore_nt_acl - *false*
- root.skip_write_permission_check - *false*
- protocol.v40_features.referrals_enabled - *false*
- protocol.v41_features.referrals_enabled - *false*
- protocol.v4_fsid_change - *true*
- protocol.v4_lease_seconds - 30
- protocol.v4_grace_seconds - 45
- protocol.v4_session_slots - 180
- protocol.v4_session_slot_reply_cache_size - 640

**Related ONTAP commands**

- vserver nfs create

- `export-policy access-cache config show`

#### Learn more

- DOC /protocols/nfs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | |
| root | root | |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>* online - NFS server is ready to accept client requests.<br>* offline - NFS server is not ready to accept client requests. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

**Example request**

```
{
  "access_cache_config": {
    "harvest_timeout": 3600,
    "ttl_failure": 1,
    "ttl_negative": 30,
    "ttl_positive": 60
  },
  "credential_cache": {
    "harvest_timeout": 86400000,
    "negative_ttl": 3600000,
    "positive_ttl": 3600000,
    "transient_error_ttl": 30000
  },
  "exports": {
    "name_service_lookup_protocol": "string"
  },
  "extended_groups_limit": 32,
  "file_session_io_grouping_count": 5000,
  "file_session_io_grouping_duration": 120,
  "protocol": {
    "v40_features": {
      "acl_max_aces": 500
    },
    "v41_features": {
      "implementation_domain": "string",
      "implementation_name": "string"
    },
    "v4_id_domain": "string"
  },
  "protocol_access_rules": {
    "cifs_access_type": "string",
    "nfs3_access_type": "string",
    "nfs4_access_type": "string"
  },
  "security": {
    "chown_mode": "string",
    "ntfs_unix_security": "string",
    "permitted_encryption_types": [
      "string"
    ]
  },
  "state": "string",
  "svm": {
    "name": "svm1",
```

```
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "transport": {
      "tcp_max_transfer_size": 16384
    },
    "windows": {
      "default_user": "string"
    }
  }
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of NFS Server Records |
| records | array[nfs_service] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "access_cache_config": {
        "harvest_timeout": 3600,
        "ttl_failure": 1,
        "ttl_negative": 30,
        "ttl_positive": 60
      },
      "credential_cache": {
        "harvest_timeout": 86400000,
        "negative_ttl": 3600000,
        "positive_ttl": 3600000,
        "transient_error_ttl": 30000
      },
      "exports": {
        "name_service_lookup_protocol": "string"
      },
      "extended_groups_limit": 32,
      "file_session_io_grouping_count": 5000,
      "file_session_io_grouping_duration": 120,
      "protocol": {
        "v40_features": {
          "acl_max_aces": 500
        },
        "v41_features": {
          "implementation_domain": "string",
          "implementation_name": "string"
        },
        "v4_id_domain": "string"
      },
      "protocol_access_rules": {
        "cifs_access_type": "string",
        "nfs3_access_type": "string",
        "nfs4_access_type": "string"
      },
      "security": {
        "chown_mode": "string",
        "ntfs_unix_security": "string",
        "permitted_encryption_types": [
          "string"
        ]
      },
```

```
      "state": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "transport": {
        "tcp_max_transfer_size": 16384
      },
      "windows": {
        "default_user": "string"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1534829 | The port numbers allowed are 635 (the default) and 1024 through 9999 |
| 2621516 | This operation is only supported on a data SVM |
| 2621570 | This operation is not permitted on an SVM that is configured as the destination for SVM DR. |
| 2621574 | This operation is not permitted on an SVM that is configured as the destination of a MetroCluster SVM relationship |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 3276852 | NFSv4.1 implementation name length for the SVM must be less than 256 bytes. |
| 3276872 | NFSv4 lease timeout value must not be less than 10. |

| Error Code | Description |
| --- | --- |
| 3276873 | NFSv4 lease timeout value must be between 10 and (protocol.v4_grace_seconds-1). |
| 3276874 | NFSv4 grace timeout value must not be greater than 180. |
| 3276875 | NFSv4 grace timeout value must be between (protocol.v4_lease_seconds+1) and 180. |
| 3276900 | NFSv4.1 implementation name cannot be an empty string |
| 3276916 | Vserver is not running |
| 3277038 | Cannot enable \"showmount\" feature because it requires an effective cluster version of Data ONTAP 8.3.0 or later |
| 3277048 | The port numbers allowed are 635 (the default) and 1024 through 9999 |
| 3277049 | Cannot enable \"showmount\" feature on ID-Discard Vserver. Ensure that the Vserver is initialized and retry the command |
| 3277052 | NFSv4.x access to transitioned volumes in this Vserver could trigger conversion of non-Unicode directories to Unicode, which might impact data-serving performance. Before enabling NFSv4.x for this Vserver, refer to the Data and Configuration Transition Guide |
| 3277069 | Cannot disable TCP because the SnapDiff RPC server is in the \"on\" state |
| 3277085 | The port numbers allowed are 1024 through 9999 |
| 3277089 | Attempting to create an NFS server using 64-bits for NFSv3 FSIDs and File IDs on Vserver. Older client software might not work with 64-bit identifiers |
| 3277099 | Domain name contains invalid characters or it is too short. Allowed characters are: alphabetical characters (A-Za-z), numeric characters (0-9), minus sign (-), and the period (.). The first character must be alphabetical or numeric, last character must not be a minus sign or a period. Minimum supported length: 2 characters, maximum of 256 characters |
| 3277140 | Cannot set "transport.tcp_max_transfer_size" to a value other than multiples of 4096 |
| 2621507 | The NFS protocol is not allowed for the specified SVM. |

| Error Code | Description |
| --- | --- |
| 2621519 | SVM name is invalid. The SVM name must begin with a letter or an underscore. If the SVM is of type "sync-source", the maximum supported length is 41. Otherwise, the maximum supported length is 47. |
| 262196 | Field "access_cache_config.harvest_timeout" cannot be set in this operation. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_cache_config

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the time after which an entry is deleted from the access cache, if unused. |
| isDnsTTLEnabled | boolean | Specifies whether Dns TTL is enabled. |
| ttl_failure | integer | Specifies the time to live value for entries for which a failure was encountered, in seconds. |
| ttl_negative | integer | Specifies the time to live value of a negative access cache, in seconds. |
| ttl_positive | integer | Specifies the time to live value of a positive access cache, in seconds. |

credential_cache

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the age of the unused cached entries, in milliseconds, after which they are cleared from the cache. |
| negative_ttl | integer | Specifies the age in milliseconds, of the negative cached credentials after which they are cleared from the cache. |

| Name | Type | Description |
| --- | --- | --- |
| positive_ttl | integer | Specifies the age in milliseconds, of the positive cached credentials after which they are cleared from the cache. |
| transient_error_ttl | integer | Specifies the age in milliseconds, of the cached entries during a transient error situation. |

exports

| Name | Type | Description |
| --- | --- | --- |
| name_service_lookup_protocol | string | Specifies the protocol to use for doing name service lookups. |
| netgroup_trust_any_nsswitch_no_match | boolean | Specifies if you can consider a no-match result from any of the netgroup ns-switch sources to be authoritative. If this option is enabled, then a no-match response from any of the netgroup ns-switch sources is deemed conclusive even if other sources could not be searched. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
|---|---|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
| --- | --- | --- |
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

v3_features

| Name | Type | Description |
| --- | --- | --- |
| connection_drop | boolean | Specifies whether the dropping of a connection when an NFSv3 request is dropped is enabled. |
| ejukebox_enabled | boolean | Specifies whether NFSv3 EJUKEBOX error is enabled. |
| fsid_change | boolean | Specifies whether the change in FSID as NFSv3 clients traverse filesystems should be shown. |
| hide_snapshot_enabled | boolean | Specifies whether hiding a snapshot directory under a NFSv3 mount point is enabled. |
| mount_daemon_port | integer | Specifies which port the NFS mount daemon (mountd) uses. |
| mount_root_only | boolean | Specifies whether the SVM allows MOUNT protocol calls only from privileged ports (port numbers less than 1024). |
| network_lock_manager_port | integer | Specifies which port the Network lock manager uses. |
| network_status_monitor_port | integer | Specifies which port the Network status monitor port uses. |
| rquota_daemon_port | integer | Specifies which port the NFS quota daemon port uses. |

v40_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.0 ACLs is enabled. |
| acl_max_aces | integer | Specifies the maximum number of aces in a NFSv4.0 ACL. |
| acl_preserve | boolean | Specifies if the NFSv4 ACL is preserved or dropped when chmod is performed. In unified security style, this parameter also specifies if NTFS file permissions are preserved or dropped when chmod, chgrp, or chown are performed. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.0 Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.0 referrals is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.0 Write Delegation is enabled. |

v41_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.1 or later ACLs is enabled. |
| implementation_domain | string | Specifies the NFSv4.1 or later implementation ID domain. |
| implementation_name | string | Specifies the NFSv4.1 or later implementation ID name. |
| pnfs_enabled | boolean | Specifies whether NFSv4.1 or later Parallel NFS is enabled. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.1 referrals is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| trunking_enabled | boolean | Specifies whether NFSv4.1 or later trunking is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Write Delegation is enabled. |

v42_features

| Name | Type | Description |
| --- | --- | --- |
| seclabel_enabled | boolean | Specifies whether NFSv4.2 or later security label is enabled. |
| sparsefile_ops_enabled | boolean | Specifies whether NFSv4.2 or later sparsefile operation is enabled. |
| xattrs_enabled | boolean | Specifies whether NFSv4.2 or later extended attributes is enabled. |

protocol

| Name | Type | Description |
| --- | --- | --- |
| v3_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled. |
| v3_enabled | boolean | Specifies whether NFSv3 protocol is enabled. |
| v3_features | v3_features | |
| v40_enabled | boolean | Specifies whether NFSv4.0 protocol is enabled. |
| v40_features | v40_features | |
| v41_enabled | boolean | Specifies whether NFSv4.1 or later protocol is enabled. |
| v41_features | v41_features | |
| v42_features | v42_features | |
| v4_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| v4_fsid_change | boolean | Specifies whether the change in FSID when NFSv4 clients traverse file systems is displayed. |
| v4_grace_seconds | integer | Specifies the grace period for clients to reclaim file locks after a server failure. |
| v4_id_domain | string | Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol. |
| v4_lease_seconds | integer | Specifies the lease seconds of the NFSv4 clients. If it is inactive for more than the time displayed, all of the file lock states on a node might be lost. |
| v4_session_slot_reply_cache_size | integer | Specifies the number of bytes of the reply that is cached in each NFSv4.x session slot. |
| v4_session_slots | integer | Specifies the number of entries in NFSv4.x session slot table. |
| v4_subnet_filter_enabled | boolean | Specifies whether NFSv4 subnet filtering is enabled. |

protocol_access_rules

| Name | Type | Description |
| --- | --- | --- |
| cifs_access_type | string | Access available for the CIFS protocol. |
| nfs3_access_type | string | Access available for the NFSv3 protocol. |
| nfs4_access_type | string | Access available for the NFSv4 protocol. |

qtree

| Name | Type | Description |
| --- | --- | --- |
| export_enabled | boolean | Specifies whether qtree export is enabled. |
| validate_export | boolean | Specifies whether qtree export validation is enabled. |

root

| Name | Type | Description |
| --- | --- | --- |
| ignore_nt_acl | boolean | Specifies whether Windows ACLs affect root access from NFS. If this option is enabled, root access from NFS ignores the NT ACL set on the file or directory. |
| skip_write_permission_check | boolean | Specifies if permission checks are to be skipped for NFS WRITE calls from root/owner. For copying read-only files to a destination folder which has inheritable ACLs, this option must be enabled. |

security

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies whether file ownership can be changed only by the superuser, or if a non-root user can also change file ownership. If you set this parameter to restricted, file ownership can be changed only by the superuser, even though the on-disk permissions allow a non-root user to change file ownership. If you set this parameter to unrestricted, file ownership can be changed by the superuser and by the non-root user, depending upon the access granted by on-disk permissions. If you set this parameter to use_export_policy, file ownership can be changed in accordance with the relevant export rules. |

| Name | Type | Description |
| --- | --- | --- |
| nt_acl_display_permission | boolean | Controls the permissions that are displayed to NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set. When true, the displayed permissions are based on the maximum access granted by the NT ACL to any user. When false, the displayed permissions are based on the minimum access granted by the NT ACL to any user. |
| ntfs_unix_security | string | Specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to ignore, ONTAP ignores NFSv3 security changes. If you set this parameter to fail, this overrides the UNIX security options set in the relevant export rules. If you set this parameter to use_export_policy, ONTAP processes NFSv3 security changes in accordance with the relevant export rules.<br>• Default value: 1<br>• enum: ["ignore", "fail", "use_export_policy"]<br>• Introduced in: 9.11<br>• x-nullable: true |
| permitted_encryption_types | array[string] | Specifies the permitted encryption types for Kerberos over NFS. |
| rpcsec_context_idle | integer | Specifies, in seconds, the amount of time a RPCSEC_GSS context is permitted to remain unused before it is deleted. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
|---|---|---|
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

transport

| Name | Type | Description |
|---|---|---|
| rdma_enabled | boolean | Specifies whether RDMA transport is enabled for the NFS server. |
| tcp_enabled | boolean | Specifies whether TCP transports are enabled on the server. |
| tcp_max_transfer_size | integer | Specifies the maximum transfer size in bytes, that the storage system negotiates with the client for TCP transport of data for NFSv3 and NFSv4.x protocols. The range is 8192 to 1048576. |
| udp_enabled | boolean | Specifies whether UDP transports are enabled on the server. |

windows

| Name | Type | Description |
| --- | --- | --- |
| default_user | string | Specifies the default Windows user for the NFS server. |
| map_unknown_uid_to_default_user | boolean | Specifies whether or not the mapping of an unknown UID to the default Windows user is enabled. |
| v3_ms_dos_client_enabled | boolean | Specifies whether NFSv3 MS-DOS client support is enabled. |

nfs_service

| Name | Type | Description |
| --- | --- | --- |
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |

| Name | Type | Description |
| --- | --- | --- |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>* online - NFS server is ready to accept client requests.<br>* offline - NFS server is not ready to accept client requests. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

nfs_service

| Name | Type | Description |
| --- | --- | --- |
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |

| Name | Type | Description |
| --- | --- | --- |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>• online - NFS server is ready to accept client requests.<br>• offline - NFS server is not ready to accept client requests. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete NFS configuration for an SVM

DELETE `/protocols/nfs/services/{svm.uuid}`

**Introduced In:** 9.6

Deletes the NFS configuration of an SVM.

#### Related ONTAP commands

- `vserver nfs delete`

#### Learn more

- DOC /protocols/nfs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 3276916 | Vserver is not running |
| 3277008 | NFS Kerberos must be disabled on all LIFs of Vserver before deleting the NFS configuration. When all LIFs are disabled, try the operation |
| 3277009 | NFS Kerberos realms associated with the Vserver are deleted |
| 3277111 | Internal error. Failed to remove NFS-specific security trace filter for Vserver |
| 3277112 | Internal error. Failed to modify the protocols field of a security trace filter for Vserver |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve NFS configuration for an SVM

GET `/protocols/nfs/services/{svm.uuid}`

**Introduced In:** 9.6

Retrieves the NFS configuration of an SVM.

#### Related ONTAP commands

- `vserver nfs show`
- `vserver nfs status`

#### Learn more

- DOC /protocols/nfs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| metric | metric | Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |

| Name | Type | Description |
| --- | --- | --- |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>• online - NFS server is ready to accept client requests.<br>• offline - NFS server is not ready to accept client requests. |
| statistics | statistics | Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_cache_config": {
    "harvest_timeout": 3600,
    "ttl_failure": 1,
    "ttl_negative": 30,
    "ttl_positive": 60
  },
  "credential_cache": {
    "harvest_timeout": 86400000,
    "negative_ttl": 3600000,
    "positive_ttl": 3600000,
    "transient_error_ttl": 30000
  },
  "exports": {
    "name_service_lookup_protocol": "string"
  },
  "extended_groups_limit": 32,
  "file_session_io_grouping_count": 5000,
  "file_session_io_grouping_duration": 120,
  "metric": {
    "v3": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
```

```
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "v4": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "v41": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
```

```
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  },
  "protocol": {
    "v40_features": {
      "acl_max_aces": 500
    },
    "v41_features": {
      "implementation_domain": "string",
      "implementation_name": "string"
    },
    "v4_id_domain": "string"
  },
  "protocol_access_rules": {
    "cifs_access_type": "string",
    "nfs3_access_type": "string",
    "nfs4_access_type": "string"
  },
  "security": {
    "chown_mode": "string",
    "ntfs_unix_security": "string",
    "permitted_encryption_types": [
      "string"
    ]
  },
  "state": "string",
  "statistics": {
    "v3": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
```

```
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "v4": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "v41": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
```

```
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "transport": {
    "tcp_max_transfer_size": 16384
  },
  "windows": {
    "default_user": "string"
  }
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_cache_config

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the time after which an entry is deleted from the access cache, if unused. |
| isDnsTTLEnabled | boolean | Specifies whether Dns TTL is enabled. |
| ttl_failure | integer | Specifies the time to live value for entries for which a failure was encountered, in seconds. |
| ttl_negative | integer | Specifies the time to live value of a negative access cache, in seconds. |
| ttl_positive | integer | Specifies the time to live value of a positive access cache, in seconds. |

credential_cache

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the age of the unused cached entries, in milliseconds, after which they are cleared from the cache. |
| negative_ttl | integer | Specifies the age in milliseconds, of the negative cached credentials after which they are cleared from the cache. |

| Name | Type | Description |
| --- | --- | --- |
| positive_ttl | integer | Specifies the age in milliseconds, of the positive cached credentials after which they are cleared from the cache. |
| transient_error_ttl | integer | Specifies the age in milliseconds, of the cached entries during a transient error situation. |

exports

| Name | Type | Description |
| --- | --- | --- |
| name_service_lookup_protocol | string | Specifies the protocol to use for doing name service lookups. |
| netgroup_trust_any_nsswitch_no_match | boolean | Specifies if you can consider a no-match result from any of the netgroup ns-switch sources to be authoritative. If this option is enabled, then a no-match response from any of the netgroup ns-switch sources is deemed conclusive even if other sources could not be searched. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
|---|---|---|
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

v3_features

| Name | Type | Description |
|---|---|---|
| connection_drop | boolean | Specifies whether the dropping of a connection when an NFSv3 request is dropped is enabled. |
| ejukebox_enabled | boolean | Specifies whether NFSv3 EJUKEBOX error is enabled. |
| fsid_change | boolean | Specifies whether the change in FSID as NFSv3 clients traverse filesystems should be shown. |
| hide_snapshot_enabled | boolean | Specifies whether hiding a snapshot directory under a NFSv3 mount point is enabled. |
| mount_daemon_port | integer | Specifies which port the NFS mount daemon (mountd) uses. |
| mount_root_only | boolean | Specifies whether the SVM allows MOUNT protocol calls only from privileged ports (port numbers less than 1024). |
| network_lock_manager_port | integer | Specifies which port the Network lock manager uses. |
| network_status_monitor_port | integer | Specifies which port the Network status monitor port uses. |
| rquota_daemon_port | integer | Specifies which port the NFS quota daemon port uses. |

v40_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.0 ACLs is enabled. |
| acl_max_aces | integer | Specifies the maximum number of aces in a NFSv4.0 ACL. |
| acl_preserve | boolean | Specifies if the NFSv4 ACL is preserved or dropped when chmod is performed. In unified security style, this parameter also specifies if NTFS file permissions are preserved or dropped when chmod, chgrp, or chown are performed. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.0 Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.0 referrals is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.0 Write Delegation is enabled. |

v41_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.1 or later ACLs is enabled. |
| implementation_domain | string | Specifies the NFSv4.1 or later implementation ID domain. |
| implementation_name | string | Specifies the NFSv4.1 or later implementation ID name. |
| pnfs_enabled | boolean | Specifies whether NFSv4.1 or later Parallel NFS is enabled. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.1 referrals is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| trunking_enabled | boolean | Specifies whether NFSv4.1 or later trunking is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Write Delegation is enabled. |

v42_features

| Name | Type | Description |
| --- | --- | --- |
| seclabel_enabled | boolean | Specifies whether NFSv4.2 or later security label is enabled. |
| sparsefile_ops_enabled | boolean | Specifies whether NFSv4.2 or later sparsefile operation is enabled. |
| xattrs_enabled | boolean | Specifies whether NFSv4.2 or later extended attributes is enabled. |

protocol

| Name | Type | Description |
| --- | --- | --- |
| v3_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled. |
| v3_enabled | boolean | Specifies whether NFSv3 protocol is enabled. |
| v3_features | v3_features | |
| v40_enabled | boolean | Specifies whether NFSv4.0 protocol is enabled. |
| v40_features | v40_features | |
| v41_enabled | boolean | Specifies whether NFSv4.1 or later protocol is enabled. |
| v41_features | v41_features | |
| v42_features | v42_features | |
| v4_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| v4_fsid_change | boolean | Specifies whether the change in FSID when NFSv4 clients traverse file systems is displayed. |
| v4_grace_seconds | integer | Specifies the grace period for clients to reclaim file locks after a server failure. |
| v4_id_domain | string | Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol. |
| v4_lease_seconds | integer | Specifies the lease seconds of the NFSv4 clients. If it is inactive for more than the time displayed, all of the file lock states on a node might be lost. |
| v4_session_slot_reply_cache_size | integer | Specifies the number of bytes of the reply that is cached in each NFSv4.x session slot. |
| v4_session_slots | integer | Specifies the number of entries in NFSv4.x session slot table. |
| v4_subnet_filter_enabled | boolean | Specifies whether NFSv4 subnet filtering is enabled. |

protocol_access_rules

| Name | Type | Description |
| --- | --- | --- |
| cifs_access_type | string | Access available for the CIFS protocol. |
| nfs3_access_type | string | Access available for the NFSv3 protocol. |
| nfs4_access_type | string | Access available for the NFSv4 protocol. |

qtree

| Name | Type | Description |
| --- | --- | --- |
| export_enabled | boolean | Specifies whether qtree export is enabled. |
| validate_export | boolean | Specifies whether qtree export validation is enabled. |

root

| Name | Type | Description |
| --- | --- | --- |
| ignore_nt_acl | boolean | Specifies whether Windows ACLs affect root access from NFS. If this option is enabled, root access from NFS ignores the NT ACL set on the file or directory. |
| skip_write_permission_check | boolean | Specifies if permission checks are to be skipped for NFS WRITE calls from root/owner. For copying read-only files to a destination folder which has inheritable ACLs, this option must be enabled. |

security

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies whether file ownership can be changed only by the superuser, or if a non-root user can also change file ownership. If you set this parameter to restricted, file ownership can be changed only by the superuser, even though the on-disk permissions allow a non-root user to change file ownership. If you set this parameter to unrestricted, file ownership can be changed by the superuser and by the non-root user, depending upon the access granted by on-disk permissions. If you set this parameter to use_export_policy, file ownership can be changed in accordance with the relevant export rules. |

| Name | Type | Description |
| --- | --- | --- |
| nt_acl_display_permission | boolean | Controls the permissions that are displayed to NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set. When true, the displayed permissions are based on the maximum access granted by the NT ACL to any user. When false, the displayed permissions are based on the minimum access granted by the NT ACL to any user. |
| ntfs_unix_security | string | Specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to ignore, ONTAP ignores NFSv3 security changes. If you set this parameter to fail, this overrides the UNIX security options set in the relevant export rules. If you set this parameter to use_export_policy, ONTAP processes NFSv3 security changes in accordance with the relevant export rules.<br>• Default value: 1<br>• enum: ["ignore", "fail", "use_export_policy"]<br>• Introduced in: 9.11<br>• x-nullable: true |
| permitted_encryption_types | array[string] | Specifies the permitted encryption types for Kerberos over NFS. |
| rpcsec_context_idle | integer | Specifies, in seconds, the amount of time a RPCSEC_GSS context is permitted to remain unused before it is deleted. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
|---|---|---|
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

transport

| Name | Type | Description |
|---|---|---|
| rdma_enabled | boolean | Specifies whether RDMA transport is enabled for the NFS server. |
| tcp_enabled | boolean | Specifies whether TCP transports are enabled on the server. |
| tcp_max_transfer_size | integer | Specifies the maximum transfer size in bytes, that the storage system negotiates with the client for TCP transport of data for NFSv3 and NFSv4.x protocols. The range is 8192 to 1048576. |
| udp_enabled | boolean | Specifies whether UDP transports are enabled on the server. |

windows

| Name | Type | Description |
|---|---|---|
| default_user | string | Specifies the default Windows user for the NFS server. |
| map_unknown_uid_to_default_user | boolean | Specifies whether or not the mapping of an unknown UID to the default Windows user is enabled. |
| v3_ms_dos_client_enabled | boolean | Specifies whether NFSv3 MS-DOS client support is enabled. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update NFS configuration for an SVM

PATCH `/protocols/nfs/services/{svm.uuid}`

**Introduced In:** 9.6

Updates the NFS configuration of an SVM.

#### Related ONTAP commands

- `vserver nfs modify`
- `vserver nfs on`
- `vserver nfs off`

- `vserver nfs start`
- `vserver nfs stop`

#### Learn more

- DOC /protocols/nfs/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |

| Name | Type | Description |
| --- | --- | --- |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>• online - NFS server is ready to accept client requests.<br>• offline - NFS server is not ready to accept client requests. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

**Example request**

```
{
  "access_cache_config": {
    "harvest_timeout": 3600,
    "ttl_failure": 1,
    "ttl_negative": 30,
    "ttl_positive": 60
  },
  "credential_cache": {
    "harvest_timeout": 86400000,
    "negative_ttl": 3600000,
    "positive_ttl": 3600000,
    "transient_error_ttl": 30000
  },
  "exports": {
    "name_service_lookup_protocol": "string"
  },
  "extended_groups_limit": 32,
  "file_session_io_grouping_count": 5000,
  "file_session_io_grouping_duration": 120,
  "protocol": {
    "v40_features": {
      "acl_max_aces": 500
    },
    "v41_features": {
      "implementation_domain": "string",
      "implementation_name": "string"
    },
    "v4_id_domain": "string"
  },
  "protocol_access_rules": {
    "cifs_access_type": "string",
    "nfs3_access_type": "string",
    "nfs4_access_type": "string"
  },
  "security": {
    "chown_mode": "string",
    "ntfs_unix_security": "string",
    "permitted_encryption_types": [
      "string"
    ]
  },
  "state": "string",
  "svm": {
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "transport": {
    "tcp_max_transfer_size": 16384
  },
  "windows": {
    "default_user": "string"
  }
}
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1534829 | The port numbers allowed are 635 (the default) and 1024 through 9999 |
| 2621516 | This operation is only supported on a data SVM |
| 2621574 | This operation is not permitted on a SVM that is configured as the destination of a MetroCluster SVM relationship |
| 3276852 | NFSv4.1 implementation name length for the SVM must be less than 256 bytes. |
| 3276872 | NFSv4 lease timeout value must not be less than 10. |
| 3276873 | NFSv4 lease timeout value must be between 10 and (protocol.v4_grace_seconds-1). |
| 3276874 | NFSv4 grace timeout value must not be greater than 180. |
| 3276875 | NFSv4 grace timeout value must be between (protocol.v4_lease_seconds+1) and 180. |
| 3276916 | Vserver is not running |
| 3276969 | Ports cannot be modified when NFS access is enabled. Disable NFS access before changing port settings |

| Error Code | Description |
| --- | --- |
| 3276994 | Kerberos must be disabled on all LIFs on the SVM before adding or removing any encryption. Disable Kerberos on the LIF and retry the operation. |
| 3277048 | The port numbers allowed are 635 (the default) and 1024 through 9999 |
| 3277069 | Cannot disable TCP because the SnapDiff RPC server is in the \"on\" state |
| 3277085 | The port numbers allowed are 1024 through 9999. |
| 3277087 | Attempting to reduce the number of bits used for NFSv3 FSIDs and File IDs from 64 to 32 on Vserver. This could result in collisions between different File IDs and is not recommended |
| 3277088 | Attempting to increase the number of bits used for NFSv3 FSIDs and File IDs from 32 to 64 on Vserver. This could result in older client software no longer working with the volumes owned by Vserver |
| 3277090 | Attempting to disallow multiple FSIDs per mount point on Vserver. Since this Vserver currently uses 32-bit NFSv3 FSIDs and File IDs, this could result in collisions between different File IDs and is not recommended |
| 3277099 | Domain name contains invalid characters or its too short. Allowed characters are: alphabetical characters (A-Za-z), numeric characters (0-9), minus sign (-), and the period (.). The first character must be alphabetical or numeric, last character must not be a minus sign or a period. Minimum supported length: 2 characters, maximum of 256 characters |
| 3277140 | Cannot set "transport.tcp_max_transfer_size" to a value other than multiples of 4096 |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_cache_config

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the time after which an entry is deleted from the access cache, if unused. |
| isDnsTTLEnabled | boolean | Specifies whether Dns TTL is enabled. |
| ttl_failure | integer | Specifies the time to live value for entries for which a failure was encountered, in seconds. |
| ttl_negative | integer | Specifies the time to live value of a negative access cache, in seconds. |
| ttl_positive | integer | Specifies the time to live value of a positive access cache, in seconds. |

credential_cache

| Name | Type | Description |
|---|---|---|
| harvest_timeout | integer | Specifies the age of the unused cached entries, in milliseconds, after which they are cleared from the cache. |
| negative_ttl | integer | Specifies the age in milliseconds, of the negative cached credentials after which they are cleared from the cache. |

| Name | Type | Description |
| --- | --- | --- |
| positive_ttl | integer | Specifies the age in milliseconds, of the positive cached credentials after which they are cleared from the cache. |
| transient_error_ttl | integer | Specifies the age in milliseconds, of the cached entries during a transient error situation. |

exports

| Name | Type | Description |
| --- | --- | --- |
| name_service_lookup_protocol | string | Specifies the protocol to use for doing name service lookups. |
| netgroup_trust_any_nsswitch_no_match | boolean | Specifies if you can consider a no-match result from any of the netgroup ns-switch sources to be authoritative. If this option is enabled, then a no-match response from any of the netgroup ns-switch sources is deemed conclusive even if other sources could not be searched. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
|---|---|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
| --- | --- | --- |
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

v3_features

| Name | Type | Description |
| --- | --- | --- |
| connection_drop | boolean | Specifies whether the dropping of a connection when an NFSv3 request is dropped is enabled. |
| ejukebox_enabled | boolean | Specifies whether NFSv3 EJUKEBOX error is enabled. |
| fsid_change | boolean | Specifies whether the change in FSID as NFSv3 clients traverse filesystems should be shown. |
| hide_snapshot_enabled | boolean | Specifies whether hiding a snapshot directory under a NFSv3 mount point is enabled. |
| mount_daemon_port | integer | Specifies which port the NFS mount daemon (mountd) uses. |
| mount_root_only | boolean | Specifies whether the SVM allows MOUNT protocol calls only from privileged ports (port numbers less than 1024). |
| network_lock_manager_port | integer | Specifies which port the Network lock manager uses. |
| network_status_monitor_port | integer | Specifies which port the Network status monitor port uses. |
| rquota_daemon_port | integer | Specifies which port the NFS quota daemon port uses. |

v40_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.0 ACLs is enabled. |
| acl_max_aces | integer | Specifies the maximum number of aces in a NFSv4.0 ACL. |
| acl_preserve | boolean | Specifies if the NFSv4 ACL is preserved or dropped when chmod is performed. In unified security style, this parameter also specifies if NTFS file permissions are preserved or dropped when chmod, chgrp, or chown are performed. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.0 Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.0 referrals is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.0 Write Delegation is enabled. |

v41_features

| Name | Type | Description |
| --- | --- | --- |
| acl_enabled | boolean | Specifies whether NFSv4.1 or later ACLs is enabled. |
| implementation_domain | string | Specifies the NFSv4.1 or later implementation ID domain. |
| implementation_name | string | Specifies the NFSv4.1 or later implementation ID name. |
| pnfs_enabled | boolean | Specifies whether NFSv4.1 or later Parallel NFS is enabled. |
| read_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Read Delegation is enabled. |
| referrals_enabled | boolean | Specifies whether NFSv4.1 referrals is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| trunking_enabled | boolean | Specifies whether NFSv4.1 or later trunking is enabled. |
| write_delegation_enabled | boolean | Specifies whether NFSv4.1 or later Write Delegation is enabled. |

v42_features

| Name | Type | Description |
| --- | --- | --- |
| seclabel_enabled | boolean | Specifies whether NFSv4.2 or later security label is enabled. |
| sparsefile_ops_enabled | boolean | Specifies whether NFSv4.2 or later sparsefile operation is enabled. |
| xattrs_enabled | boolean | Specifies whether NFSv4.2 or later extended attributes is enabled. |

protocol

| Name | Type | Description |
| --- | --- | --- |
| v3_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled. |
| v3_enabled | boolean | Specifies whether NFSv3 protocol is enabled. |
| v3_features | v3_features | |
| v40_enabled | boolean | Specifies whether NFSv4.0 protocol is enabled. |
| v40_features | v40_features | |
| v41_enabled | boolean | Specifies whether NFSv4.1 or later protocol is enabled. |
| v41_features | v41_features | |
| v42_features | v42_features | |
| v4_64bit_identifiers_enabled | boolean | Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| v4_fsid_change | boolean | Specifies whether the change in FSID when NFSv4 clients traverse file systems is displayed. |
| v4_grace_seconds | integer | Specifies the grace period for clients to reclaim file locks after a server failure. |
| v4_id_domain | string | Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol. |
| v4_lease_seconds | integer | Specifies the lease seconds of the NFSv4 clients. If it is inactive for more than the time displayed, all of the file lock states on a node might be lost. |
| v4_session_slot_reply_cache_size | integer | Specifies the number of bytes of the reply that is cached in each NFSv4.x session slot. |
| v4_session_slots | integer | Specifies the number of entries in NFSv4.x session slot table. |
| v4_subnet_filter_enabled | boolean | Specifies whether NFSv4 subnet filtering is enabled. |

protocol_access_rules

| Name | Type | Description |
| --- | --- | --- |
| cifs_access_type | string | Access available for the CIFS protocol. |
| nfs3_access_type | string | Access available for the NFSv3 protocol. |
| nfs4_access_type | string | Access available for the NFSv4 protocol. |

qtree

| Name | Type | Description |
| --- | --- | --- |
| export_enabled | boolean | Specifies whether qtree export is enabled. |
| validate_export | boolean | Specifies whether qtree export validation is enabled. |

root

| Name | Type | Description |
| --- | --- | --- |
| ignore_nt_acl | boolean | Specifies whether Windows ACLs affect root access from NFS. If this option is enabled, root access from NFS ignores the NT ACL set on the file or directory. |
| skip_write_permission_check | boolean | Specifies if permission checks are to be skipped for NFS WRITE calls from root/owner. For copying read-only files to a destination folder which has inheritable ACLs, this option must be enabled. |

security

| Name | Type | Description |
| --- | --- | --- |
| chown_mode | string | Specifies whether file ownership can be changed only by the superuser, or if a non-root user can also change file ownership. If you set this parameter to restricted, file ownership can be changed only by the superuser, even though the on-disk permissions allow a non-root user to change file ownership. If you set this parameter to unrestricted, file ownership can be changed by the superuser and by the non-root user, depending upon the access granted by on-disk permissions. If you set this parameter to use_export_policy, file ownership can be changed in accordance with the relevant export rules. |

| Name | Type | Description |
| --- | --- | --- |
| nt_acl_display_permission | boolean | Controls the permissions that are displayed to NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set. When true, the displayed permissions are based on the maximum access granted by the NT ACL to any user. When false, the displayed permissions are based on the minimum access granted by the NT ACL to any user. |
| ntfs_unix_security | string | Specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to ignore, ONTAP ignores NFSv3 security changes. If you set this parameter to fail, this overrides the UNIX security options set in the relevant export rules. If you set this parameter to use_export_policy, ONTAP processes NFSv3 security changes in accordance with the relevant export rules.<br>• Default value: 1<br>• enum: ["ignore", "fail", "use_export_policy"]<br>• Introduced in: 9.11<br>• x-nullable: true |
| permitted_encryption_types | array[string] | Specifies the permitted encryption types for Kerberos over NFS. |
| rpcsec_context_idle | integer | Specifies, in seconds, the amount of time a RPCSEC_GSS context is permitted to remain unused before it is deleted. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
|---|---|---|
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

transport

| Name | Type | Description |
|---|---|---|
| rdma_enabled | boolean | Specifies whether RDMA transport is enabled for the NFS server. |
| tcp_enabled | boolean | Specifies whether TCP transports are enabled on the server. |
| tcp_max_transfer_size | integer | Specifies the maximum transfer size in bytes, that the storage system negotiates with the client for TCP transport of data for NFSv3 and NFSv4.x protocols. The range is 8192 to 1048576. |
| udp_enabled | boolean | Specifies whether UDP transports are enabled on the server. |

windows

| Name | Type | Description |
| --- | --- | --- |
| default_user | string | Specifies the default Windows user for the NFS server. |
| map_unknown_uid_to_default_user | boolean | Specifies whether or not the mapping of an unknown UID to the default Windows user is enabled. |
| v3_ms_dos_client_enabled | boolean | Specifies whether NFSv3 MS-DOS client support is enabled. |

nfs_service

| Name | Type | Description |
| --- | --- | --- |
| access_cache_config | access_cache_config | |
| auth_sys_extended_groups_enabled | boolean | Specifies whether or not extended groups support over AUTH_SYS is enabled. |
| credential_cache | credential_cache | |
| enabled | boolean | Specifies if the NFS service is administratively enabled. |
| exports | exports | |
| extended_groups_limit | integer | Specifies the maximum auxiliary groups supported over AUTH_SYS and RPCSEC_GSS. |
| file_session_io_grouping_count | integer | Number of I/O operations on a file to be grouped and considered as one session for event generation applications, such as FPolicy. |
| file_session_io_grouping_duration | integer | The duration for which I/O operations on a file will be grouped and considered as one session for event generation applications, such as FPolicy. |
| protocol | protocol | |
| protocol_access_rules | protocol_access_rules | |
| qtree | qtree | |
| root | root | |

| Name | Type | Description |
| --- | --- | --- |
| rquota_enabled | boolean | Specifies whether or not the remote quota feature is enabled. |
| security | security | |
| showmount_enabled | boolean | Specifies whether or not the showmount feature is enabled. |
| state | string | Specifies the state of the NFS service on the SVM. The following values are supported:<br>• online - NFS server is ready to accept client requests.<br>• offline - NFS server is not ready to accept client requests. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport | transport | |
| vstorage_enabled | boolean | Specifies whether or not the VMware vstorage feature is enabled. |
| windows | windows | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Retrieve NFS protocol historical performance metrics

GET `/protocols/nfs/services/{svm.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for the NFS protocol of an SVM.

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| v4.status | string | query | False | Filter by v4.status<br>• Introduced in: 9.8 |
| v4.duration | string | query | False | Filter by v4.duration<br>• Introduced in: 9.8 |
| v4.throughput.write | integer | query | False | Filter by v4.throughput.write<br>• Introduced in: 9.8 |
| v4.throughput.read | integer | query | False | Filter by v4.throughput.read<br>• Introduced in: 9.8 |
| v4.throughput.total | integer | query | False | Filter by v4.throughput.total<br>• Introduced in: 9.8 |
| v4.latency.other | integer | query | False | Filter by v4.latency.other<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| v4.latency.total | integer | query | False | Filter by v4.latency.total<br><br>• Introduced in: 9.8 |
| v4.latency.write | integer | query | False | Filter by v4.latency.write<br><br>• Introduced in: 9.8 |
| v4.latency.read | integer | query | False | Filter by v4.latency.read<br><br>• Introduced in: 9.8 |
| v4.iops.other | integer | query | False | Filter by v4.iops.other<br><br>• Introduced in: 9.8 |
| v4.iops.total | integer | query | False | Filter by v4.iops.total<br><br>• Introduced in: 9.8 |
| v4.iops.write | integer | query | False | Filter by v4.iops.write<br><br>• Introduced in: 9.8 |
| v4.iops.read | integer | query | False | Filter by v4.iops.read<br><br>• Introduced in: 9.8 |
| timestamp | string | query | False | Filter by timestamp |
| v41.throughput.write | integer | query | False | Filter by v41.throughput.write<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| v41.throughput.read | integer | query | False | Filter by v41.throughput.read<br><br>• Introduced in: 9.8 |
| v41.throughput.total | integer | query | False | Filter by v41.throughput.total<br><br>• Introduced in: 9.8 |
| v41.iops.other | integer | query | False | Filter by v41.iops.other<br><br>• Introduced in: 9.8 |
| v41.iops.total | integer | query | False | Filter by v41.iops.total<br><br>• Introduced in: 9.8 |
| v41.iops.write | integer | query | False | Filter by v41.iops.write<br><br>• Introduced in: 9.8 |
| v41.iops.read | integer | query | False | Filter by v41.iops.read<br><br>• Introduced in: 9.8 |
| v41.latency.other | integer | query | False | Filter by v41.latency.other<br><br>• Introduced in: 9.8 |
| v41.latency.total | integer | query | False | Filter by v41.latency.total<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| v41.latency.write | integer | query | False | Filter by v41.latency.write<br><br>• Introduced in: 9.8 |
| v41.latency.read | integer | query | False | Filter by v41.latency.read<br><br>• Introduced in: 9.8 |
| v41.duration | string | query | False | Filter by v41.duration<br><br>• Introduced in: 9.8 |
| v41.status | string | query | False | Filter by v41.status<br><br>• Introduced in: 9.8 |
| v3.iops.other | integer | query | False | Filter by v3.iops.other |
| v3.iops.total | integer | query | False | Filter by v3.iops.total |
| v3.iops.write | integer | query | False | Filter by v3.iops.write |
| v3.iops.read | integer | query | False | Filter by v3.iops.read |
| v3.latency.other | integer | query | False | Filter by v3.latency.other |
| v3.latency.total | integer | query | False | Filter by v3.latency.total |
| v3.latency.write | integer | query | False | Filter by v3.latency.write |
| v3.latency.read | integer | query | False | Filter by v3.latency.read |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| v3.throughput.write | integer | query | False | Filter by v3.throughput.write |
| v3.throughput.read | integer | query | False | Filter by v3.throughput.read |
| v3.throughput.total | integer | query | False | Filter by v3.throughput.total |
| v3.duration | string | query | False | Filter by v3.duration |
| v3.status | string | query | False | Filter by v3.status |
| svm.uuid | string | path | True | Unique identifier of the SVM. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "v3": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "v4": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
```

```
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "v41": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      }
    }
  ]
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

v3

The NFSv3 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v4

The NFSv4 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

v41

The NFSv4.1 operations

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

records

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

| Name | Type | Description |
|---|---|---|
| v3 | v3 | The NFSv3 operations |
| v4 | v4 | The NFSv4 operations |
| v41 | v41 | The NFSv4.1 operations |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NFS over TLS interface configuration details

### Manage NFS over TLS interface configuration details

#### Examples

### Retrieving the NFS over TLS interface configuration details

```
# The API:
GET /api/protocols/nfs/tls/interfaces

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/tls/interfaces"
```

### Updating the NFS over TLS interface configuration

```
# The API:
PATCH /api/protocols/nfs/tls/interfaces/{interface.uuid}

# The call:
curl -d "@test_patch_nfs_tls_interface.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/tls/interfaces/e62936de-7342-11e8-9eb4-0050568be2b7"
test_patch_nfs_tls_interface.txt(body):
{
 "enabled" : "true"
}
```

### Retrieve NFS over TLS interface configurations

GET `/protocols/nfs/tls/interfaces`

**Introduced In:** 9.15

Retrieves NFS over TLS interfaces.

#### Related ONTAP commands

- `vserver nfs tls interface show`

#### Learn more

- DOC /protocols/nfs/tls/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.ip.address | string | query | False | Filter by interface.ip.address |
| interface.name | string | query | False | Filter by interface.name |
| interface.uuid | string | query | False | Filter by interface.uuid |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| certificate.name | string | query | False | Filter by certificate.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| certificate.uuid | string | query | False | Filter by certificate.uuid |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[nfs_tls_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

certificate

Specifies the certificate that is used for creating NFS over TLS connections.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Network interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nfs_tls_interface

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| certificate | certificate | Specifies the certificate that is used for creating NFS over TLS connections. |
| enabled | boolean | Indicates whether NFS over TLS is enabled. |
| interface | interface | Network interface. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NFS over TLS interface configuration

GET `/protocols/nfs/tls/interfaces/{interface.uuid}`

**Introduced In:** 9.15

Retrieves an NFS over TLS interface.

#### Related ONTAP commands

- `vserver nfs tls interface show`

#### Learn more

- DOC /protocols/nfs/tls/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.uuid | string | path | True | Network interface UUID. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| certificate | certificate | Specifies the certificate that is used for creating NFS over TLS connections. |
| enabled | boolean | Indicates whether NFS over TLS is enabled. |
| interface | interface | Network interface. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
| --- | --- |
| 262197 | The value provided is invalid for field "fields". |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

certificate

Specifies the certificate that is used for creating NFS over TLS connections.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Network interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the properties of an NFS over TLS interface**

PATCH `/protocols/nfs/tls/interfaces/{interface.uuid}`

**Introduced In:** 9.15

Updates the properties of an NFS over TLS interface.

#### Optional query parameter

- `skip-san-validation` - Specifies whether the server should ignore validating the certificate for Subject Alternate Name.

#### Related ONTAP commands

- `vserver nfs tls interface modify`
- `vserver nfs tls interface enable`
- `vserver nfs tls interface disable`

#### Learn more

- DOC /protocols/nfs/tls/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.uuid | string | path | True | Network interface UUID. |
| skip_san_validation | boolean | query | False | Indicates whether Subject Alternate Name (SAN) validation has to be skipped.<br>• Introduced in: 9.17 |

#### Request Body

| Name | Type | Description |
|---|---|---|
| certificate | certificate | Specifies the certificate that is used for creating NFS over TLS connections. |
| enabled | boolean | Indicates whether NFS over TLS is enabled. |
| interface | interface | Network interface. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "interface": {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error codes | Description |
|---|---|
| 2 | The value provided is an invalid value for field "interface.uuid". |
| 4 | The entry doesn’t exist |
| 3277204 | TLS already enabled/disabled on this LIF. |
| 3277205 | Cannot enable TLS because no certificate was provided. |
| 3277206 | The "enabled" is a required field. |
| 3277210 | The FQDN of the LIF IP address and the common name present in the certificate do not match. Update the certificate so that the common name and LIF FQDN match and retry the operation. |

| Error codes | Description |
| --- | --- |
| 3277217 | Failed to enable TLS on LIF on vserver because the certificate does not have LIF as subject alternate name. |
| 92405900 | Certificate not found for the SVM. |
| 92406020 | Only certificates of type 'server' are supported on the SVM. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

certificate

Specifies the certificate that is used for creating NFS over TLS connections.

| Name | Type | Description |
|---|---|---|
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ip

IP information

interface

Network interface.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nfs_tls_interface

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Specifies the certificate that is used for creating NFS over TLS connections. |
| enabled | boolean | Indicates whether NFS over TLS is enabled. |
| interface | interface | Network interface. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and create Vscan configuration

## View and create Vscan configuration

### Overview

Use Vscan to protect data from being compromised by viruses or other malicious code. Vscan combines best-in-class third party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from third party vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software. Vscan is not supported on continuous availability (CA) shares.

### Examples

## Retrieving all of the Vscan configurations

```
# The API:
/api/protocols/vscan

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/vscan?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "03ce5c36-f269-11e8-8852-0050568e5298",
      "name": "vs1"
    },
    "enabled": true,
    "scanner_pools": [
      {
        "name": "scanner-1",
        "servers": [
          "1.1.1.1",
          "10.72.204.27"
        ],
        "privileged_users": [
          "cifs\\u1",
          "cifs\\u2"
        ],
        "role": "primary",
        "cluster": {
          "name": "Cluster1",
          "uuid": "0228714d-f268-11e8-8851-0050568e5298"
        }
```

```
      },
      {
        "name": "scanner-2",
        "servers": [
          "1.1.1.1",
          "10.72.204.27"
        ],
        "privileged_users": [
          "cifs\\u1",
          "cifs\\u2"
        ],
        "role": "primary",
        "cluster": {
          "name": "Cluster1",
          "uuid": "0228714d-f268-11e8-8851-0050568e5298"
        }
      }
    ],
    "on_access_policies": [
      {
        "name": "default_CIFS",
        "vsName": "vs1",
        "enabled": true,
        "mandatory": true,
        "scope": {
          "max_file_size": 2147483648,
          "include_extensions": [
            "*"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        },
        "protocol": "CIFS",
        "owner": "cluster"
      },
      {
        "name": "on-access-test1",
        "vsName": "vs1",
        "enabled": false,
        "mandatory": true,
        "scope": {
          "max_file_size": 10000,
          "exclude_paths": [
            "\\dir"
          ],
```

```
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        },
        "protocol": "CIFS",
        "owner": "vserver"
      },
      {
        "name": "on-access-test2",
        "vsName": "vs1",
        "enabled": false,
        "mandatory": true,
        "scope": {
          "max_file_size": 10000,
          "exclude_paths": [
            "\\dir"
          ],
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        },
        "protocol": "CIFS",
        "owner": "vserver"
      }
    ],
    "on_demand_policies": [
      {
        "name": "task-1",
        "scan_paths": [
          "/vol1"
```

```
        ],
        "log_path": "/vol1",
        "scope": {
          "max_file_size": 10000,
          "exclude_paths": [
            "/vol1"
          ],
          "include_extensions": [
            "vmdk",
            "mp*"
          ],
          "exclude_extensions": [
            "mp3",
            "mp4"
          ],
          "scan_without_extension": true
        }
      },
      {
        "name": "task-2",
        "scan_paths": [
          "/vol1"
        ],
        "log_path": "/vol2",
        "scope": {
          "max_file_size": 10000,
          "exclude_paths": [
            "/vol2"
          ],
          "include_extensions": [
            "vmdk",
            "mp*"
          ],
          "exclude_extensions": [
            "mp3",
            "mp4"
          ],
          "scan_without_extension": true
        }
      }
    ]
  },
  {
    "svm": {
      "uuid": "24c2567a-f269-11e8-8852-0050568e5298",
      "name": "vs2"
```

```
    },
    "enabled": false,
    "scanner_pools": [
      {
        "name": "sp2",
        "servers": [
          "1.1.1.1"
        ],
        "privileged_users": [
          "cifs\\u1"
        ],
        "role": "idle"
      }
    ],
    "on_access_policies": [
      {
        "name": "default_CIFS",
        "vsName": "vs2",
        "enabled": true,
        "mandatory": true,
        "scope": {
          "max_file_size": 2147483648,
          "include_extensions": [
            "*"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        }
      },
      {
        "name": "ap1",
        "vsName": "vs2",
        "enabled": false,
        "mandatory": true,
        "scope": {
          "max_file_size": 2147483648,
          "include_extensions": [
            "*"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        }
      }
    ],
```

```
    "on_demand_policies": [
      {
        "name": "t1",
        "scan_paths": [
          "/vol1"
        ],
        "log_path": "/vol1",
        "scope": {
          "max_file_size": 10737418240,
          "include_extensions": [
            "*"
          ],
          "scan_without_extension": true
        }
      }
    ]
  }
],
"num_records": 2
}
```

**Retrieving all Vscan configurations for a particular SVM**

```
# The API:
/api/protocols/vscan/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/24c2567a-f269-11e8-
8852-0050568e5298?fields=*" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "24c2567a-f269-11e8-8852-0050568e5298",
  "name": "vs2"
},
"enabled": false,
"scanner_pools": [
  {
    "name": "sp2",
    "servers": [
      "1.1.1.1"
    ],
    "privileged_users": [
      "cifs\\u1"
```

```
    ],
    "role": "idle"
  }
],
"on_access_policies": [
  {
    "name": "default_CIFS",
    "vsName": "vs2",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    },
    "protocol": "CIFS",
    "owner": "cluster"
  },
  {
    "name": "ap1",
    "vsName": "vs2",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    },
    "protocol": "CIFS",
    "owner": "vserver"
  }
],
"on_demand_policies": [
  {
    "name": "t1",
    "scan_paths": [
      "/vol1"
    ],
```

```
    "log_path": "/vol1",
    "scope": {
      "max_file_size": 10737418240,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true
    }
  }
]
}
```

#### Creating a Vscan configuration

```
# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"enabled\": true, \"on_access_policies\": [ { \"enabled\": true,
\"mandatory\": true, \"name\": \"on-access-test\", \"scope\": {
\"exclude_extensions\": [ \"mp*\", \"txt\" ], \"exclude_paths\": [
\"\\\\vol\" ], \"include_extensions\": [ \"mp*\", \"txt\" ],
\"max_file_size\": 21474, \"only_execute_access\": false,
\"scan_readonly_volumes\": false, \"scan_without_extension\": true } } ],
\"on_demand_policies\": [ { \"log_path\": \"/vol\", \"name\": \"task-1\",
\"scan_paths\": [ \"/vol\" ], \"schedule\": { \"name\": \"daily\",
\"uuid\": \"d4984822-17b7-11e9-b450-0050568ecd85\" }, \"scope\": {
\"exclude_extensions\": [ \"mp3\", \"mp4\" ], \"exclude_paths\": [
\"/vol\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
\"max_file_size\": 10737, \"scan_without_extension\": true } } ],
\"scanner_pools\": [ { \"cluster\": { \"name\": \"Cluster1\", \"uuid\":
\"ab746d77-17b7-11e9-b450-0050568ecd85\" }, \"name\": \"scanner-1\",
\"privileged_users\": [ \"cifs\\\\u1\", \"cifs\\\\u2\" ], \"role\":
\"primary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ] } ], \"svm\":
{ \"name\": \"vs1\", \"uuid\": \"b103be27-17b8-11e9-b451-0050568ecd85\"
}}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
```

```
    "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
    "name": "vs1"
    },
    "enabled": true,
    "scanner_pools": [
      {
        "name": "scanner-1",
        "servers": [
          "1.1.1.1",
          "10.72.204.27"
        ],
        "privileged_users": [
          "cifs\\u1",
          "cifs\\u2"
        ],
        "role": "primary",
        "cluster": {
          "name": "Cluster1",
          "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
        }
      }
    ],
    "on_access_policies": [
      {
        "name": "on-access-test",
        "enabled": true,
        "mandatory": true,
        "scope": {
          "max_file_size": 21474,
          "exclude_paths": [
            "\\vol"
          ],
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        }
      }
    ],
```

```
    "on_demand_policies": [
      {
        "name": "task-1",
        "scan_paths": [
          "/vol"
        ],
        "log_path": "/vol",
        "schedule": {
          "uuid": "d4984822-17b7-11e9-b450-0050568ecd85",
          "name": "daily"
        },
        "scope": {
          "max_file_size": 10737,
          "exclude_paths": [
            "//"
          ],
          "include_extensions": [
            "vmdk",
            "mp*"
          ],
          "exclude_extensions": [
            "mp3",
            "mp4"
          ],
          "scan_without_extension": true
        }
      }
    ]
  }
]
}
```


**Creating multiple Vscan scanner-pools for the specified SVM**

```
# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"scanner_pools\": [ { \"cluster\": { \"name\": \"Cluster1\", \"uuid\":
\"ab746d77-17b7-11e9-b450-0050568ecd85\" }, \"name\": \"scanner-1\",
\"privileged_users\": [ \"cifs\\\\u1\", \"cifs\\\\u2\" ], \"role\":
\"primary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ] }, {
\"cluster\": { \"name\": \"Cluster1\", \"uuid\": \"ab746d77-17b7-11e9-
```

```
b450-0050568ecd85\" }, \"name\": \"scanner-2\", \"privileged_users\": [
\"cifs\\\\u3\", \"cifs\\\\u4\" ], \"role\": \"primary\", \"servers\": [
\"1.1.1.5\", \"10.72.3.27\" ] } ], \"svm\": { \"name\": \"vs1\", \"uuid\":
\"b103be27-17b8-11e9-b451-0050568ecd85\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
    "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
    "name": "vs1"
    },
    "scanner_pools": [
      {
        "name": "scanner-1",
        "servers": [
          "1.1.1.1",
          "10.72.204.27"
        ],
        "privileged_users": [
          "cifs\\u1",
          "cifs\\u2"
        ],
        "role": "primary",
        "cluster": {
          "name": "Cluster1",
          "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
        }
      },
      {
        "name": "scanner-2",
        "servers": [
          "1.1.1.5",
          "10.72.3.27"
        ],
        "privileged_users": [
          "cifs\\u3",
          "cifs\\u4"
        ],
        "role": "primary",
        "cluster": {
          "name": "Cluster1",
          "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
        }
```

```
      }
    ]
  }
]
}
```

#### Creating multiple Vscan On-access policies for a specified SVM

```
# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"on_access_policies\": [ { \"enabled\": false, \"mandatory\": true,
\"name\": \"on-access-test11\", \"scope\": { \"exclude_extensions\": [
\"mp*\", \"txt\" ], \"exclude_paths\": [ \"\\\\vol\" ],
\"include_extensions\": [ \"mp*\", \"txt\" ], \"max_file_size\": 214748,
\"only_execute_access\": false, \"scan_readonly_volumes\": false,
\"scan_without_extension\": true } }, { \"enabled\": false, \"mandatory\":
true, \"name\": \"on-access-test10\", \"scope\": { \"exclude_extensions\":
[ \"mp*\", \"txt\" ], \"exclude_paths\": [ \"\\\\vol\" ],
\"include_extensions\": [ \"mp*\", \"txt\" ], \"max_file_size\": 21474,
\"only_execute_access\": false, \"scan_readonly_volumes\": false,
\"scan_without_extension\": true } } ], \"svm\": { \"name\": \"vs1\",
\"uuid\": \"b103be27-17b8-11e9-b451-0050568ecd85\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
      "name": "vs1"
    },
    "on_access_policies": [
      {
        "name": "on-access-test11",
        "enabled": false,
        "mandatory": true,
        "scope": {
          "max_file_size": 214748,
          "exclude_paths": [
            "\\vol"
```

```
          ],
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        }
      },
      {
        "name": "on-access-test10",
        "enabled": false,
        "mandatory": true,
        "scope": {
          "max_file_size": 21474,
          "exclude_paths": [
            "\\vol"
          ],
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "scan_without_extension": true,
          "scan_readonly_volumes": false,
          "only_execute_access": false
        }
      }
    ]
  }
]
}
```


**Creating multiple Vscan On-demand policies for a specified SVM**

```
# The API:
/api/protocols/vscan
```

```
# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"on_demand_policies\": [ { \"log_path\": \"/vol\", \"name\": \"task-1\",
\"scan_paths\": [ \"/vol\" ], \"schedule\": { \"name\": \"daily\",
\"uuid\": \"d4984822-17b7-11e9-b450-0050568ecd85\" }, \"scope\": {
\"exclude_extensions\": [ \"mp3\", \"mp4\" ], \"exclude_paths\": [
\"/vol1\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
\"max_file_size\": 107374, \"scan_without_extension\": true } }, {
\"log_path\": \"/vol\", \"name\": \"task-2\", \"scan_paths\": [ \"/vol\"
], \"scope\": { \"exclude_extensions\": [ \"mp3\", \"mp4\" ],
\"exclude_paths\": [ \"/vol1\" ], \"include_extensions\": [ \"vmdk\",
\"mp*\" ], \"max_file_size\": 107374, \"scan_without_extension\": true } }
], \"svm\": { \"name\": \"vs1\", \"uuid\": \"b103be27-17b8-11e9-b451-
0050568ecd85\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
      "name": "vs1"
    },
    "on_demand_policies": [
      {
        "name": "task-1",
        "scan_paths": [
          "/vol"
        ],
        "log_path": "/vol",
        "schedule": {
          "uuid": "d4984822-17b7-11e9-b450-0050568ecd85",
          "name": "daily"
        },
        "scope": {
          "max_file_size": 107374,
          "exclude_paths": [
            "/vol1"
          ],
          "include_extensions": [
            "vmdk",
            "mp*"
          ],
```

```
          "exclude_extensions": [
            "mp3",
            "mp4"
          ],
          "scan_without_extension": true
        }
      },
      {
        "name": "task-2",
        "scan_paths": [
          "/vol"
        ],
        "log_path": "/vol",
        "scope": {
          "max_file_size": 107374,
          "exclude_paths": [
            "/vol1"
          ],
          "include_extensions": [
            "vmdk",
            "mp*"
          ],
          "exclude_extensions": [
            "mp3",
            "mp4"
          ],
          "scan_without_extension": true
        }
      }
    ]
  }
]
}
```

#### Enabling Vscan for a specified SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-
8852-0050568e5298" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"enabled\": true}"
```

#### Clearing the Vscan cache for the specified SVM

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-
8852-0050568e5298" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"cache_clear\": true}"
```

#### Deleting the Vscan configuration for a specified SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-
8852-0050568e5298" -H "accept: application/json"
```

### Retrieve the Vscan configuration

GET `/protocols/vscan`

**Introduced In:** 9.6

Retrieves the Vscan configuration.
This includes scanner-pools, On-Access policies, On-Demand policies, and information about whether a Vscan is enabled or disabled on an SVM.

Important notes:

- You can enable only one Vscan configuration at a time for an SVM.
- You can only query using `svm.uuid` or `svm.name`.

#### Related ONTAP commands

- `vserver vscan show`
- `vserver vscan scanner-pool show`
- `vserver vscan scanner-pool servers show`
- `vserver vscan scanner-pool privileged-users show`
- `vserver vscan on-access-policy show`
- `vserver vscan on-access-policy file-ext-to-exclude show`
- `vserver vscan on-access-policy file-ext-to-include show`
- `vserver vscan on-access-policy paths-to-exclude show`
- `vserver vscan on-demand-task show`

**Learn more**

- DOC /protocols/vscan
- DOC /protocols/vscan/{svm.uuid}/scanner-pools

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scanner_pools.role | string | query | False | Filter by scanner_pools.role |
| scanner_pools.name | string | query | False | Filter by scanner_pools.name<br>• maxLength: 256<br>• minLength: 1 |
| scanner_pools.cluster.name | string | query | False | Filter by scanner_pools.cluster.name |
| scanner_pools.cluster.uuid | string | query | False | Filter by scanner_pools.cluster.uuid |
| scanner_pools.privileged_users | string | query | False | Filter by scanner_pools.privileged_users |
| scanner_pools.servers | string | query | False | Filter by scanner_pools.servers |
| on_demand_policies.scan_paths | string | query | False | Filter by on_demand_policies.scan_paths |
| on_demand_policies.scope.exclude_extensions | string | query | False | Filter by on_demand_policies.scope.exclude_extensions |
| on_demand_policies.scope.include_extensions | string | query | False | Filter by on_demand_policies.scope.include_extensions |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| on_demand_policies.scope.max_file_size | integer | query | False | Filter by on_demand_policies.scope.max_file_size<br>• Max value: 1099511627776<br>• Min value: 1024 |
| on_demand_policies.scope.exclude_paths | string | query | False | Filter by on_demand_policies.scope.exclude_paths<br>• maxLength: 255<br>• minLength: 1 |
| on_demand_policies.scope.scan_without_extension | boolean | query | False | Filter by on_demand_policies.scope.scan_without_extension |
| on_demand_policies.log_path | string | query | False | Filter by on_demand_policies.log_path |
| on_demand_policies.schedule.name | string | query | False | Filter by on_demand_policies.schedule.name |
| on_demand_policies.schedule.uuid | string | query | False | Filter by on_demand_policies.schedule.uuid |
| on_demand_policies.name | string | query | False | Filter by on_demand_policies.name<br>• maxLength: 256<br>• minLength: 1 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| on_access_policies.name | string | query | False | Filter by on_access_policies.name<br>• maxLength: 256<br>• minLength: 1 |
| on_access_policies.scope.max_file_size | integer | query | False | Filter by on_access_policies.scope.max_file_size<br>• Max value: 1099511627776<br>• Min value: 1024 |
| on_access_policies.scope.only_execute_access | boolean | query | False | Filter by on_access_policies.scope.only_execute_access |
| on_access_policies.scope.include_extensions | string | query | False | Filter by on_access_policies.scope.include_extensions |
| on_access_policies.scope.scan_readonly_volumes | boolean | query | False | Filter by on_access_policies.scope.scan_readonly_volumes |
| on_access_policies.scope.exclude_extensions | string | query | False | Filter by on_access_policies.scope.exclude_extensions |
| on_access_policies.scope.scan_without_extension | boolean | query | False | Filter by on_access_policies.scope.scan_without_extension |
| on_access_policies.scope.exclude_paths | string | query | False | Filter by on_access_policies.scope.exclude_paths<br>• maxLength: 255<br>• minLength: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| on_access_policies.mandatory | boolean | query | False | Filter by on_access_policies.mandatory |
| on_access_policies.enabled | boolean | query | False | Filter by on_access_policies.enabled |
| on_access_policies.protocol | string | query | False | Filter by on_access_policies.protocol<br>• Introduced in: 9.18 |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[vscan] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "on_access_policies": [
        {
          "name": "on-access-test",
          "protocol": "string",
          "scope": {
            "exclude_extensions": [
              "mp*",
              "txt"
            ],
            "exclude_paths": [
              "\\dir1\\dir2\\name",
              "\\vol\\a b",
              "\\vol\\a,b\\"
            ],
            "include_extensions": [
              "mp*",
              "txt"
            ],
            "max_file_size": 2147483648
          }
        }
      ],
      "on_demand_policies": [
        {
          "log_path": "/vol0/report_dir",
          "name": "task-1",
          "scan_paths": [
```

```
            "/vol1/",
            "/vol2/cifs/"
          ],
          "schedule": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "scope": {
            "exclude_extensions": [
              "mp3",
              "mp4"
            ],
            "exclude_paths": [
              "/vol1/cold-files/",
              "/vol1/cifs/names"
            ],
            "include_extensions": [
              "vmdk",
              "mp*"
            ],
            "max_file_size": 10737418240
          }
        }
      ],
      "scanner_pools": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "name": "scanner-1",
          "privileged_users": [
            "cifs\\u1",
            "cifs\\u2"
          ],
          "role": "string",
```

```
            "servers": [
              "1.1.1.1",
              "10.72.204.27",
              "vmwin204-27.fsct.nb"
            ]
          }
        ],
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_access_policy

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for

viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |

schedule

Schedule of the task.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |

| Name | Type | Description |
|---|---|---|
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_demand_policy

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |

cluster_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |

| Name | Type | Description |
| --- | --- | --- |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.10 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.10 |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |
| on_access_policies | array[vscan_on_access_policy] | |
| on_demand_policies | array[vscan_on_demand_policy] | |
| scanner_pools | array[scanner_pool] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a Vscan configuration

POST `/protocols/vscan`

**Introduced In:** 9.6

Creates a Vscan configuration, which includes a list of scanner-pools, Vscan On-Access policies and Vscan On-Demand policies. Defines whether the Vscan configuration you create is enabled or disabled for a specified SVM.

Important notes:

- You can enable only one Vscan configuration at a time for an SVM.
- There needs to be at least one active scanner-pool and one enabled On-Access policy to enable Vscan successfully.
- By default, a Vscan is enabled when it's created.
- By default, the Vscan On-Access policies created from this endpoint are in the disabled state. You can use the On-Access policy PATCH endpoint to enable a particular On-Access policy. In ONTAP 9.6, only one Vscan On-Access policy can be enabled and only one Vscan On-Demand policy can be scheduled on an SVM.
- Vscan is not supported on continuous availability (CA) shares.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the Vscan configuration.

#### Recommended optional properties

- `scanner_pools` - There must be at least one active scanner-pool for Vscan configuration. Created either through Vscan POST operation or scanner-pools POST operation.

#### Default property values

If not specified in POST, the following default property value is assigned:

- `enabled` - *true*

#### Related ONTAP commands

- `vserver vscan enable`

- `vserver vscan scanner-pool create`
- `vserver vscan scanner-pool apply-policy`
- `vserver vscan scanner-pool servers add`
- `vserver vscan scanner-pool privileged-users add`
- `vserver vscan on-access-policy create`
- `vserver vscan on-access-policy file-ext-to-exclude add`
- `vserver vscan on-access-policy file-ext-to-include add`
- `vserver vscan on-access-policy paths-to-exclude add`
- `vserver vscan on-demand-task create`

#### Learn more

- DOC /protocols/vscan
- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |
| on_access_policies | array[vscan_on_access_policy] | |
| on_demand_policies | array[vscan_on_demand_policy] | |
| scanner_pools | array[scanner_pool] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "on_access_policies": [
    {
      "name": "on-access-test",
      "protocol": "string",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": 2147483648
      }
    }
  ],
  "on_demand_policies": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
```

```
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "max_file_size": 10737418240
      }
    }
  ],
  "scanner_pools": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "string",
      "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
      ]
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[vscan] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "on_access_policies": [
        {
          "name": "on-access-test",
          "protocol": "string",
          "scope": {
            "exclude_extensions": [
              "mp*",
              "txt"
            ],
            "exclude_paths": [
              "\\dir1\\dir2\\name",
              "\\vol\\a b",
              "\\vol\\a,b\\"
            ],
            "include_extensions": [
              "mp*",
              "txt"
            ],
            "max_file_size": 2147483648
          }
        }
      ],
      "on_demand_policies": [
        {
          "log_path": "/vol0/report_dir",
          "name": "task-1",
          "scan_paths": [
            "/vol1/",
            "/vol2/cifs/"
          ],
          "schedule": {
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "scope": {
            "exclude_extensions": [
              "mp3",
              "mp4"
            ],
            "exclude_paths": [
```

```
                "/vol1/cold-files/",
                "/vol1/cifs/names"
              ],
              "include_extensions": [
                "vmdk",
                "mp*"
              ],
              "max_file_size": 10737418240
            }
          }
        ],
        "scanner_pools": [
          {
            "cluster": {
              "name": "cluster1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "name": "scanner-1",
            "privileged_users": [
              "cifs\\u1",
              "cifs\\u2"
            ],
            "role": "string",
            "servers": [
              "1.1.1.1",
              "10.72.204.27",
              "vmwin204-27.fsct.nb"
            ]
          }
        ],
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027259 + | A scanner-pool, an On-Access policy, or an On-Demand policy might fail to get created due to either a systematic error or some hardware failure. The error code returned details the failure along with the reason for the failure. For example, if a scanner-pool fails due to an incorrect cluster name, then the error might read: "Failed to create scanner-pool "scanner-1". Reason: "Cluster uuid points to different cluster name instead of the cluster-name supplied.". Retry the operation." |
| 10027260 + | If a scanner-pool, an On-Access policy or an On-Demand policy specified in the input already exists, then a duplicate error is returned. For example, if a scanner-pool "scanner-1" already exists for an SVM and is again specified in the input, the error message will read: " Failed to create scanner-pool "scanner-1" as the specified entry already exists. Delete the entry and retry the POST operation." |
| 2621462 + | The specified SVM name is invalid |
| 2621706 + | The specified svm.uuid is either invalid or belongs to a different SVM |
| 10027015 + | Attempting to enable a Vscan but no active scanner-pool exists for the specified SVM |
| 10027011 + | Attempting to enable a Vscan for an SVM for which no CIFS server exists |
| 10027023 + | Attempting to enable a Vscan for an SVM for which no active Vscan On-Access policy exist |
| 10027086 + | DNS resolution failed for one or more hostnames |
| 10027012 + | Cannot enable Vscan on an administrative SVM. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_access_policy

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |

| Name | Type | Description |
| --- | --- | --- |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |

schedule

Schedule of the task.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
| --- | --- | --- |
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_demand_policy

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the

scope of an On-Demand scan.

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |

cluster_reference

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |

| Name | Type | Description |
| --- | --- | --- |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.10 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.10 |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

| Name | Type | Description |
| --- | --- | --- |
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |
| on_access_policies | array[vscan_on_access_policy] | |
| on_demand_policies | array[vscan_on_demand_policy] | |
| scanner_pools | array[scanner_pool] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage Vscan configuration

### Manage Vscan configuration

#### Overview

This API is used to display connection status information for the external virus-scanning servers or "Vscan servers".

#### Examples

### Retrieving all fields for the Vscan server status

```
# The API:
/api/protocols/vscan/server-status/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/server-status?fields=*"
-H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "66f8564d-aefc-11eb-bd8c-0050568e8ed1",
      "name": "vs1"
    },
    "node": {
      "uuid": "ce2463d9-aef6-11eb-bd8c-0050568e8ed1",
      "name": "user-vsim1"
    },
    "ip": "10.140.132.141",
    "type": "primary",
    "state": "connected",
    "update_time": "2021-05-07T21:35:02+05:30",
    "vendor": "mighty master anti-evil scanner",
    "version": "5.643",
```

```
    "extended_stats": [
      "ts=10:15:28 PM May 07,2021",
      "scans=sent:0,compok:0,comperr:0,compnotfnd:0,ms/comp:0",
      "avshim-version=1.0.5.0",
      "mempage/s=0, procs=52, threads=2063, %cpu=0.22, procqlen=0,
diskio/s=0, smbbytes/s=52, ",
      "ifmac=FA:16:3E:F0:70:88,
tcpstat=retrans:7538,connfail:2137,connreset:52018,inerr:0",
      "cfg=Host Name:SCSPB2271817001",
      "OS Name:Microsoft Windows Server 2012 R2 Standard",
      "OS Version:6.3.9600 N/A Build 9600",
      "System Boot Time:5/6/2021, 9:01:42 PM",
      "System Manufacturer:RDO",
      "System Model:OpenStack Compute",
      "System Type:x64-based PC",
      "Processor(s):2 Processor(s) Installed.",
      "[01]:Intel64 Family 6 Model 94 Stepping 3 GenuineIntel ~2400 Mhz"
    ],
    "interface": {
      "name": "vs1.data",
      "uuid": "6911e7c6-aefc-11eb-bd8c-0050568e8ed1",
      "ip": {
        "address": "10.140.69.165"
      }
    }
  },
  {
    "svm": {
      "uuid": "a776e8f2-aef9-11eb-8530-0050568e8ed1",
      "name": "vs2"
    },
    "node": {
      "uuid": "ce2463d9-aef6-11eb-bd8c-0050568e8ed1",
      "name": "user-vsim1"
    },
    "ip": "10.140.128.163",
    "type": "primary",
    "state": "connected",
    "update_time": "2021-05-07T21:35:43+05:30",
    "vendor": "mighty master anti-evil scanner",
    "version": "5.643",
    "extended_stats": [
      "ts=10:15:26 PM May 07,2021",
      "scans=sent:0,compok:0,comperr:0,compnotfnd:0,ms/comp:0",
      "avshim-version=1.0.5.0",
      "mempage/s=0, procs=46, threads=1947, %cpu=0.07, procqlen=0,
```

```
diskio/s=1, smbbytes/s=181, ",
      "ifmac=FA:16:3E:06:9F:D6,
tcpstat=retrans:123453,connfail:6864,connreset:13188,inerr:0",
      "cfg=Host Name:SCSPB2271815001",
      "OS Name:Microsoft Windows Server 2012 R2 Standard",
      "OS Version:6.3.9600 N/A Build 9600",
      "System Boot Time:4/20/2021, 5:23:29 PM",
      "System Manufacturer:RDO",
      "System Model:OpenStack Compute",
      "System Type:x64-based PC",
      "Processor(s):2 Processor(s) Installed.",
      "[01]:Intel64 Family 6 Model 94 Stepping 3 GenuineIntel ~2400 Mhz"
    ],
    "interface": {
      "name": "vs2.data",
      "uuid": "c070b4c2-aef9-11eb-8530-0050568e8ed1",
      "ip": {
        "address": "10.140.70.154"
      }
    }
  }
],
"num_records": 2
}
```


#### Retrieving the server status information for the server with IP address 10.141.46.173

```
# The API:
/api/protocols/vscan/server-status

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/server-
status?ip=10.140.132.141&fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "66f8564d-aefc-11eb-bd8c-0050568e8ed1",
      "name": "vs1"
    },
```

```
    "node": {
      "uuid": "ce2463d9-aef6-11eb-bd8c-0050568e8ed1",
      "name": "user-vsim1"
    },
    "ip": "10.140.132.141",
    "type": "primary",
    "state": "connected",
    "update_time": "2021-05-07T23:08:21+05:30",
    "vendor": "mighty master anti-evil scanner",
    "version": "5.643",
    "extended_stats": [
      "ts=12:07:30 AM May 08,2021",
      "scans=sent:0,compok:0,comperr:0,compnotfnd:0,ms/comp:0",
      "avshim-version=1.0.5.0",
      "mempage/s=0, procs=52, threads=2108, %cpu=0.07, procqlen=0,
diskio/s=0, smbbytes/s=230, ",
      "ifmac=FA:16:3E:F0:70:88,
tcpstat=retrans:8249,connfail:2174,connreset:52243,inerr:0",
      "cfg=Host Name:SCSPB2271817001",
      "OS Name:Microsoft Windows Server 2012 R2 Standard",
      "OS Version:6.3.9600 N/A Build 9600",
      "System Boot Time:5/6/2021, 9:01:42 PM",
      "System Manufacturer:RDO",
      "System Model:OpenStack Compute",
      "System Type:x64-based PC",
      "Processor(s):2 Processor(s) Installed.",
      "[01]:Intel64 Family 6 Model 94 Stepping 3 GenuineIntel ~2400 Mhz"
    ],
    "interface": {
      "name": "vs1.data",
      "uuid": "6911e7c6-aefc-11eb-bd8c-0050568e8ed1",
      "ip": {
        "address": "10.140.69.165"
      }
    }
  }
],
"num_records": 1
}
```


---

#### Retrieve the Vscan server status

GET `/protocols/vscan/server-status`

**Introduced In:** 9.6

Retrieves a Vscan server status.

#### Related ONTAP commands

- `vserver vscan connection-status show-all`

#### Learn more

- DOC /protocols/vscan/server-status

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| type | string | query | False | Filter by type |
| vendor | string | query | False | Filter by vendor |
| update_time | string | query | False | Filter by update_time |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| interface.ip.address | string | query | False | Filter by interface.ip.address<br>• Introduced in: 9.10 |
| interface.name | string | query | False | Filter by interface.name<br>• Introduced in: 9.10 |
| interface.uuid | string | query | False | Filter by interface.uuid<br>• Introduced in: 9.10 |
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| disconnected_reason | string | query | False | Filter by disconnected_reason |
| ip | string | query | False | Filter by ip |
| version | string | query | False | Filter by version |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[vscan_server_status] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "disconnected_reason": "string",
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "ip": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
        },
        "type": "string",
        "update_time": "string",
        "vendor": "string",
        "version": "string"
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Address of the interface used for the Vscan connection.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_server_status

Displays the connection status information of the external virus-scanning servers.

| Name | Type | Description |
| --- | --- | --- |
| disconnected_reason | string | Specifies the server disconnected reason.<br>The following is a list of the possible reasons:<br>• unknown - Disconnected, unknown reason.<br>• vscan_disabled - Disconnected, Vscan is disabled on the SVM.<br>• no_data_lif - Disconnected, SVM does not have data LIF.<br>• session_uninitialized - Disconnected, session is not initialized.<br>• remote_closed - Disconnected, server has closed the connection.<br>• invalid_protocol_msg - Disconnected, invalid protocol message received.<br>• invalid_session_id - Disconnected, invalid session ID received.<br>• inactive_connection - Disconnected, no activity on connection.<br>• invalid_user - Connection request by an invalid user.<br>• server_removed - Disconnected, server has been removed from the active Scanners List.<br>enum:<br>• unknown<br>• vscan_disabled<br>• no_data_lif<br>• session_uninitialized<br>• remote_closed<br>• invalid_protocol_msg<br>• invalid_session_id<br>• inactive_connection<br>• invalid_user<br>• server_removed |

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | Address of the interface used for the Vscan connection. |
| ip | string | IP address of the Vscan server. |
| node | node | |
| state | string | Specifies the server connection state indicating if it is in the connected or disconnected state. The following is a list of the possible states:<br>• connected - Connected<br>• disconnected - Disconnected enum:<br>• connected<br>• disconnected |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Server type. The possible values are:<br>• primary - Primary server<br>• backup - Backup server |
| update_time | string | Specifies the time the server is in the connected or disconnected state. |
| vendor | string | Name of the connected virus-scanner vendor. |
| version | string | Version of the connected virus-scanner. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a Vscan configuration

DELETE `/protocols/vscan/{svm.uuid}`

**Introduced In:** 9.6

Deletes a Vscan configuration.

Important notes:

- The Vscan DELETE endpoint deletes all of the Vscan configuration of an SVM. It first disables the Vscan and then deletes all of the SVM scanner-pools, On-Access policies, and On-Demand policies.
- Disable the active Vscan On-Access policy on an SVM before performing the Vscan delete operation on that SVM.

#### Related ONTAP commands

- `vserver vscan scanner-pool delete`
- `vserver vscan on-access-policy delete`
- `vserver vscan on-demand-policy delete`

#### Learn more

- DOC /protocols/vscan
- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027259 + | A scanner-pool, an On-Access policy, or an On-Demand policy might fail to get deleted due to either a systematic error or some hardware failure. The error code returned details the failure along with the reason for the failure. For example, "Failed to delete On-Access policy "sp1". Reason: "Failed to delete policy. Reason: policy must be disabled before being deleted.". Retry the operation." |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the Vscan configuration for an SVM

GET `/protocols/vscan/{svm.uuid}`

**Introduced In:** 9.6

Retrieves the Vscan configuration for a specified SVM.
This includes scanner-pools, On-Access policies, On-Demand policies, and information about whether a Vscan is enabled or disabled on an SVM.

Important note:

- You can enable only one Vscan configuration at a time for an SVM.

#### Related ONTAP commands

- `vserver vscan show`
- `vserver vscan scanner-pool show`
- `vserver vscan scanner-pool servers show`
- `vserver vscan scanner-pool privileged-users show`
- `vserver vscan on-access-policy show`
- `vserver vscan on-access-policy file-ext-to-exclude show`
- `vserver vscan on-access-policy file-ext-to-include show`

- `vserver vscan on-access-policy paths-to-exclude show`
- `vserver vscan on-demand-task show`

#### Learn more

- DOC /protocols/vscan
- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |
| on_access_policies | array[vscan_on_access_policy] | |
| on_demand_policies | array[vscan_on_demand_policy] | |
| scanner_pools | array[scanner_pool] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "on_access_policies": [
    {
      "name": "on-access-test",
      "protocol": "string",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": 2147483648
      }
    }
  ],
  "on_demand_policies": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "max_file_size": 10737418240
      }
    }
  ],
  "scanner_pools": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "string",
      "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
      ]
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_access_policy

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |

| Name | Type | Description |
| --- | --- | --- |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |

schedule

Schedule of the task.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
| --- | --- | --- |
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_demand_policy

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |

cluster_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |

| Name | Type | Description |
| --- | --- | --- |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.10 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.10 |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the Vscan configuration for an SVM**

PATCH `/protocols/vscan/{svm.uuid}`

**Introduced In:** 9.6

Updates the Vscan configuration of an SVM. Allows you to either enable or disable a Vscan, and allows you to clear the Vscan cache that stores the past scanning data for an SVM.

Important note:

- The Vscan PATCH endpoint does not allow you to modify scanner-pools, On-Demand policies or On-Access policies. Those modifications can only be done through their respective endpoints.

**Related ONTAP commands**

- `vserver vscan enable`
- `vserver vscan disable`
- `vserver vscan reset`

#### Learn more

- DOC /protocols/vscan
- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027015 + | Attempting to enable a Vscan but no active scanner-pool exists for the specified SVM |
| 10027009 + | Attempting to enable a Vscan for an SVM for which it's already enabled |
| 10027010 + | Attempting to disable a Vscan for an SVM for which it's already disabled |
| 10027011 + | Attempting to enable a Vscan for an SVM for which no CIFS server exists |

| Error Code | Description |
| --- | --- |
| 10027023 + | Attempting to enable a Vscan for an SVM for which no active Vscan On-Access policy exists |
| 10027012 + | Cannot enable Vscan on an administrative SVM. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_access_policy

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |

| Name | Type | Description |
| --- | --- | --- |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |

schedule

Schedule of the task.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
| --- | --- | --- |
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

vscan_on_demand_policy

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The path from the Vserver root where the task report is created. |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |

cluster_reference

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster_reference | |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.10 |

| Name | Type | Description |
| --- | --- | --- |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.10 |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

| Name | Type | Description |
|---|---|---|
| cache_clear | boolean | Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only |
| enabled | boolean | Specifies whether or not Vscan is enabled on the SVM. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve Vscan events

GET `/protocols/vscan/{svm.uuid}/events`

**Introduced In:** 9.11

Retrieves Vscan events, which are generated by the cluster to capture important events.

#### Related ONTAP commands

- `vserver vscan show-events`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scan_engine_status | integer | query | False | Filter by scan_engine_status<br>• Introduced in: 9.18 |
| type | string | query | False | Filter by type |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| vendor | string | query | False | Filter by vendor |
| server | string | query | False | Filter by server |
| consecutive_occurrence_count | integer | query | False | Filter by consecutive_occurrence_count<br>• Introduced in: 9.18 |
| event_time | string | query | False | Filter by event_time |
| svm.name | string | query | False | Filter by svm.name |
| interface.ip.address | string | query | False | Filter by interface.ip.address |
| interface.name | string | query | False | Filter by interface.name |
| interface.uuid | string | query | False | Filter by interface.uuid |
| version | string | query | False | Filter by version |
| disconnect_reason | string | query | False | Filter by disconnect_reason |
| file_path | string | query | False | Filter by file_path |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[vscan_event] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "consecutive_occurrence_count": 2,
      "disconnect_reason": "string",
      "event_time": "2021-11-25 04:29:41 +0000",
      "file_path": "/1",
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scan_engine_status": 333300005,
      "server": "192.168.1.1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "type": "string",
        "vendor": "mighty master anti-evil scanner",
        "version": "1.0"
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Address of the interface used for the Vscan connection.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_event

| Name | Type | Description |
| --- | --- | --- |
| consecutive_occurrence_count | integer | Specifies the consecutive occurrence count of the event. |

| Name | Type | Description |
| --- | --- | --- |
| disconnect_reason | string | Specifies the reason of the Vscan server disconnection.<br>The available values are:<br>• na Not applicable<br>• vscan_disabled Vscan disabled on the SVM<br>• no_data_lif SVM does not have data lif on the node<br>• session_uninitialized Session not initialized<br>• remote_closed Closure from Server<br>• invalid_protocol_msg Invalid protocol-message received<br>• invalid_session_id Invalid session-id received<br>• inactive_connection No activity on connection<br>• invalid_user Connection request by invalid user<br>• server_removed Server removed from the active scanner-pool |
| event_time | string | Specifies the Timestamp of the event. |
| file_path | string | Specifies the file for which event happened. |
| interface | interface | Address of the interface used for the Vscan connection. |
| node | node | |
| scan_engine_status | integer | Specifies the status code of the scan-engine for the event. This field is only generated when there is a proper response from the scan-engine. |
| server | string | Specifies the IP address of the Vscan server. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the event type. |
| vendor | string | Specifies the scan-engine vendor. |
| version | string | Specifies the scan-engine version. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage Vscan On-Access policies

### Manage Vscan On-Access policies

#### Overview

Use Vscan On-Access scanning to actively scan file objects for viruses when clients access files over SMB. To control which file operations trigger a vscan, use Vscan File-Operations Profile (vscan-fileop-profile) option in the CIFS share. The Vscan On-Access policy configuration defines the scope and status of On-Access scanning on file objects. Use this API to retrieve and manage Vscan On-Access policy configurations and Vscan On-Access policy statuses for the SVM.

#### Examples

### Retrieving all fields for all policies of an SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/{svm.uuid}/on-access-
policies?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "name": "default_CIFS",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    },
    "protocol": "CIFS",
    "owner": "cluster",
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-
005056b41bd1/on-access-policies/default_CIFS"
      }
    }
  },
```

```
  {
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "name": "on-access-policy",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 3221225472,
      "exclude_paths": [
        "\\vol\\a b\\",
        "\\vol\\a,b\\"
      ],
      "include_extensions": [
        "mp*",
        "tx*"
      ],
      "exclude_extensions": [
        "mp3",
        "txt"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": true
    },
    "protocol": "CIFS",
    "owner": "vserver",
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-
005056b41bd1/on-access-policies/on-access-policy"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-
access-policies?fields=*"
```

```
  }
}
}
```

#### Retrieving the specific On-Access policy associated with the specified SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/179d3c85-7053-11e8-
b9b8-005056b41bd1/on-access-policies/on-access-policy" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"name": "on-access-policy",
"enabled": true,
"mandatory": true,
"scope": {
  "max_file_size": 3221225472,
  "exclude_paths": [
    "\\vol\\a b\\",
    "\\vol\\a,b\\"
  ],
  "include_extensions": [
    "mp*",
    "tx*"
  ],
  "exclude_extensions": [
    "mp3",
    "txt"
  ],
```

```
  "scan_without_extension": true,
  "scan_readonly_volumes": false,
  "only_execute_access": true
},
"_links": {
  "self": {
    "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-
access-policies/task1"
  }
}
}
```

#### Creating a Vscan On-Access policy

The Vscan On-Access policy POST endpoint creates an On-Access policy for the specified SVM. Set enabled to "true" to enable scanning on the created policy.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"mandatory\": true, \"name\": \"on-access-policy\", \"scope\": {
\"exclude_extensions\": [ \"txt\", \"mp3\" ], \"exclude_paths\": [
\"\\\\dir1\\\\dir2\\\\ame\", \"\\\\vol\\\\a b\" ], \"include_extensions\":
[  \"mp*\", \"txt\" ], \"max_file_size\": 3221225472,
\"only_execute_access\": true, \"scan_readonly_volumes\": false,
\"scan_without_extension\": true }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "name": "on-access-policy",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 3221225472,
```

```
      "exclude_paths": [
        "\\dir1\\dir2\\ame",
        "\\vol\\a b"
      ],
      "include_extensions": [
        "mp*",
        "txt"
      ],
      "exclude_extensions": [
        "txt",
        "mp3"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": true
    }
  }
]
}
```

### Creating a Vscan On-Access policy where a number of optional fields are not specified

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"mandatory\": true, \"name\": \"on-access-policy\", \"scope\": {
\"exclude_paths\": [ \"\\\\vol\\\\a b\", \"\\\\vol\\\\a,b\\\\\" ],
\"max_file_size\": 1073741824, \"scan_without_extension\": true }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "name": "on-access-policy",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 1073741824,
      "exclude_paths": [
        "\\vol\\a b",
        "\\vol\\a,b\\"
      ],
      "scan_without_extension": true
    }
  }
]
}
```

### Updating a Vscan On-Access policy

The policy being modified is identified by the UUID of the SVM and the policy name.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies/on-access-policy" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"scope\":
{ \"include_extensions\": [ \"txt\" ], \"only_execute_access\": true,
\"scan_readonly_volumes\": false, \"scan_without_extension\": true }}"
```

---

### Deleting a Vscan On-Access policy

The policy to be deleted is identified by the UUID of the SVM and the policy name.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies/on-access-policy" -H "accept:
application/hal+json"
```

---

### Retrieve a Vscan On-Access policy

GET `/protocols/vscan/{svm.uuid}/on-access-policies`

**Introduced In:** 9.6

Retrieves the Vscan On-Access policy.

#### Related ONTAP commands

- `vserver vscan on-access-policy show`
- `vserver vscan on-access-policy file-ext-to-include show`
- `vserver vscan on-access-policy file-ext-to-exclude show`
- `vserver vscan on-access-policy paths-to-exclude show`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-access-policies

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol | string | query | False | Filter by protocol<br>• Introduced in: 9.18 |
| mandatory | boolean | query | False | Filter by mandatory |
| enabled | boolean | query | False | Filter by enabled |
| scope.scan_readonly_volumes | boolean | query | False | Filter by scope.scan_readonly_volumes |
| scope.exclude_extensions | string | query | False | Filter by scope.exclude_extensions |
| scope.max_file_size | integer | query | False | Filter by scope.max_file_size<br>• Max value: 1099511627776<br>• Min value: 1024 |
| scope.only_execute_access | boolean | query | False | Filter by scope.only_execute_access |
| scope.include_extensions | string | query | False | Filter by scope.include_extensions |
| scope.exclude_paths | string | query | False | Filter by scope.exclude_paths<br>• maxLength: 255<br>• minLength: 1 |
| scope.scan_without_extension | boolean | query | False | Filter by scope.scan_without_extension |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| owner | string | query | False | Filter by owner<br>• Introduced in: 9.18 |
| name | string | query | False | Filter by name<br>• maxLength: 256<br>• minLength: 1 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.16 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[vscan_on_access] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "name": "on-access-test",
      "owner": "string",
      "protocol": "string",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": 2147483648
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_access

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a Vscan On-Access policy

POST `/protocols/vscan/{svm.uuid}/on-access-policies`

**Introduced In:** 9.6

Creates a Vscan On-Access policy.
</b>Important notes:

- You must enable the policy on an SVM before its files can be scanned.
- You can enable only one On-Access policy at a time on an SVM. By default, the policy is enabled on creation. * If the Vscan On-Access policy has been created successfully on an SVM but cannot be enabled due to an error, the Vscan On-Access policy configurations are saved. The Vscan On-Access policy is then enabled using the PATCH operation.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the Vscan On-Access policy.
- `name` - Name of the Vscan On-Access policy. Maximum length is 256 characters.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `mandatory` - *true*
- `include_extensions` - *
- `max_file_size` - *2147483648*
- `only_execute_access` - *false*

- `scan_readonly_volumes` - *false*
- `scan_without_extension` - *true*
- `protocol` - *CIFS*

#### Related ONTAP commands

- `vserver vscan on-access-policy create`
- `vserver vscan on-access-policy enable`
- `vserver vscan on-access-policy disable`
- `vserver vscan on-access-policy file-ext-to-include add`
- `vserver vscan on-access-policy file-ext-to-exclude add`
- `vserver vscan on-access-policy paths-to-exclude add`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-access-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |

| Name | Type | Description |
|---|---|---|
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "name": "on-access-test",
  "owner": "string",
  "protocol": "string",
  "scope": {
    "exclude_extensions": [
      "mp*",
      "txt"
    ],
    "exclude_paths": [
      "\\dir1\\dir2\\name",
      "\\vol\\a b",
      "\\vol\\a,b\\"
    ],
    "include_extensions": [
      "mp*",
      "txt"
    ],
    "max_file_size": 2147483648
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[vscan_on_access] | |

#### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "name": "on-access-test",
      "owner": "string",
      "protocol": "string",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": 2147483648
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027043 | The new On-Access policy cannot be created because the SVM has reached the maximum limit of 10 On-Access policies. Delete an existing policy in order to create a new policy |
| 10027101 | The file size must be in the range 1KB to 1TB |
| 10027107 | The include extensions list cannot be empty. Specify at least one extension for inclusion |
| 10027109 | The specified CIFS path is invalid. It must be in the form "\dir1\dir2" or "\dir1\dir2\" |
| 10027249 | The On-Access policy created successfully but failed to enable the policy. The reason for enable policy operation failure might be that another policy is enabled. Disable the enabled policy and then enable the newly created policy using the PATCH operation |
| 10027253 | The number of paths specified exceeds the configured number of maximum paths. You cannot specify more than the maximum number of configured paths |
| 10027254 | The number of extensions specified exceeds the configured maximum number of extensions. You cannot specify more than the maximum number of configured extensions |

#### Definitions

**See Definitions**

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_access

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete an antivirus On-Access policy configuration

DELETE `/protocols/vscan/{svm.uuid}/on-access-policies/{name}`

**Introduced In:** 9.6

Deletes the anti-virus On-Access policy configuration.

#### Related ONTAP commands

- `vserver vscan on-access-policy delete`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-access-policies

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 10027034 | An On-Access policy associated with a data SVM which was created by SVM owned by the cluster cannot be deleted. |
| 10027040 | An On-Access policy with a status enabled cannot be deleted. Disable the policy and then delete the policy. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the Vscan On-Access policy configuration for an SVM

GET `/protocols/vscan/{svm.uuid}/on-access-policies/{name}`

**Introduced In:** 9.6

Retrieves the Vscan On-Access policy configuration of an SVM.

#### Related ONTAP commands

- `vserver vscan on-access-policy show`
- `vserver vscan on-access-policy file-ext-to-include show`
- `vserver vscan on-access-policy file-ext-to-exclude show`
- `vserver vscan on-access-policy paths-to-exclude show`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-access-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| name | string | On-Access policy name |

| Name | Type | Description |
| --- | --- | --- |
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "name": "on-access-test",
  "owner": "string",
  "protocol": "string",
  "scope": {
    "exclude_extensions": [
      "mp*",
      "txt"
    ],
    "exclude_paths": [
      "\\dir1\\dir2\\name",
      "\\vol\\a b",
      "\\vol\\a,b\\"
    ],
    "include_extensions": [
      "mp*",
      "txt"
    ],
    "max_file_size": 2147483648
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the Vscan On-Access policy configuration for an SVM

PATCH `/protocols/vscan/{svm.uuid}/on-access-policies/{name}`

**Introduced In:** 9.6

Updates the Vscan On-Access policy configuration and/or enables/disables the Vscan On-Access policy of an SVM. You cannot modify the configurations for an On-Access policy associated with a data SVM which was created by SVM owned by the cluster, although you can enable and disable the policy associated with cluster SVM.

#### Related ONTAP commands

- `vserver vscan on-access-policy modify`
- `vserver vscan on-access-policy enable`
- `vserver vscan on-access-policy disable`

- `vserver vscan on-access-policy file-ext-to-include add`
- `vserver vscan on-access-policy file-ext-to-exclude add`
- `vserver vscan on-access-policy paths-to-exclude add`
- `vserver vscan on-access-policy file-ext-to-include remove`
- `vserver vscan on-access-policy file-ext-to-exclude remove`
- `vserver vscan on-access-policy paths-to-exclude remove`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-access-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "owner": "string",
  "protocol": "string",
  "scope": {
    "exclude_extensions": [
      "mp*",
      "txt"
    ],
    "exclude_paths": [
      "\\dir1\\dir2\\name",
      "\\vol\\a b",
      "\\vol\\a,b\\"
    ],
    "include_extensions": [
      "mp*",
      "txt"
    ],
    "max_file_size": 2147483648
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027033 | Configurations for an On-Access policy associated with a data SVM which was created by SVM owned by the cluster cannot be modified. However, the policy can be enabled or disabled. |

| Error Code | Description |
| --- | --- |
| 10027046 | The specified SVM is not the owner of the specified policy. Check for the correct SVM who owns the policy. |
| 10027101 | The file size must be in the range 1KB to 1TB |
| 10027107 | The include extensions list cannot be empty. Specify at least one extension for inclusion. |
| 10027109 | The specified CIFS path is invalid. It must be in the form "\dir1\dir2" or "\dir1\dir2\". |
| 10027249 | The On-Access policy updated successfully but failed to enable/disable the policy. The reason for an enable policy operation failure might be that another policy is enabled. Disable the already enabled policy and then enable the policy. The reason for a disable policy operation failure might be that Vscan is enabled on the SVM. Disable the Vscan first and then disable the policy. |
| 10027250 | The On-Access policy cannot be enabled/disabled. The reason for an enable policy operation failure might be that another policy is enabled. Disable the already enabled policy and then enable the policy. The reason for a disable policy operation failure might be that Vscan is enabled on the SVM. Disable the Vscan and then disable the policy. |
| 10027253 | The number of paths specified exceeds the configured maximum number of paths. You cannot specify more than the maximum number of configured paths. |
| 10027254 | The number of extensions specified exceeds the configured maximum number of extensions. You cannot specify more than the maximum number of configured extensions. |

**Definitions**

**See Definitions**

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| only_execute_access | boolean | Scan only files opened with execute-access. |
| scan_readonly_volumes | boolean | Specifies whether or not read-only volume can be scanned. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_access

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Status of the On-Access Vscan policy |
| mandatory | boolean | Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning. |
| owner | string | Specifies the on-access policy owner. |
| protocol | string | Specifies the file access protocol for the on-access policy. The following lists the possible protocols. CIFS - SMB protocol |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage Vscan On-Demand policies

### Manage Vscan On-Demand policies

#### Overview

Vscan On-Demand scanning is used to check files for viruses on a schedule. For example, it can be used to run scans only in off-peak hours, or to scan very large files that are excluded from an on-access scan. Vscan On-Demand scanning can be used for any path in the SVM namespace.

Vscan On-Demand policy configurations define the scope of a Vscan On-Demand scan. The schedule parameter in the On-Demand policy configuration decides when to execute the task. Schedule can be created using the /api/clusters/schedule endpoint and can be assigned on policy create or policy modify. This API is used to retrieve and manage Vscan On-Demand policy configurations. It is also used to schedule the Vscan On-Demand scan.

#### Examples

### Retrieving all fields for all policies of an SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies/

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/{svm.uuid}/on-demand-
policies?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
        }
      }
    },
    "name": "on-demand-policy1",
```

```
    "scan_paths": [
      "/vol1/",
      "/vol2/cifs/"
    ],
    "log_path": "/vol0/report_dir",
    "schedule": {
      "uuid": "f6d0843e-f159-11e8-8e22-0050568e0945",
      "name": "schedule",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/f6d0843e-f159-11e8-8e22-
0050568e0945"
        }
      }
    },
    "scope": {
      "max_file_size": 10737418240,
      "exclude_paths": [
        "/vol1/cold-files/",
        "/vol1/cifs/names"
      ],
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "exclude_extensions": [
        "mp3",
        "mp4"
      ],
      "scan_without_extension": false
    },
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-
0050568e0945/on-demand-policies/policy1"
      }
    }
  },
  {
    "svm": {
      "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
        }
```

```
      }
    },
    "name": "on-demand-policy2",
    "scan_paths": [
      "/vol1/",
      "/vol2/cifs/"
    ],
    "log_path": "/report",
    "scope": {
      "max_file_size": 10737418240,
      "include_extensions": [
        "mp*"
      ],
      "scan_without_extension": true
    },
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-
0050568e0945/on-demand-policies/policy2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-
demand-policies?fields=*"
  }
}
}
```


### Retrieving a specific On-Demand policy associated with a specified SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-demand-policies/on-demand-task" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
    }
  }
},
"name": "on-demand-policy",
"scan_paths": [
  "/vol1/cifs"
],
"log_path": "/report",
"scope": {
  "max_file_size": 10737418240,
  "include_extensions": [
    "vmdk",
    "mp*"
  ],
  "scan_without_extension": true
},
"_links": {
  "self": {
    "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-
demand-policies/policy2"
  }
}
}
```

### Creating a Vscan On-Demand policy

The Vscan On-Demand policy POST endpoint creates an On-Demand policy for the specified SVM. Specify the schedule parameter to schedule an On-Demand scan.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-demand-policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"log_path\":
\"/vol0/report_dir\", \"name\": \"on-demand-policy\", \"scan_paths\": [
\"/vol1/\", \"/vol2/cifs/\" ], \"schedule\": { \"name\": \"weekly\",
\"uuid\": \"1cd8a442-86d1-11e0-ae1c-123478563412\" }, \"scope\": {
\"exclude_extensions\": [ \"mp3\" ], \"exclude_paths\": [ \"/vol/cold-
files/\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
\"max_file_size\": 1073741824, \"scan_without_extension\": true }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "name": "on-demand-policy",
    "scan_paths": [
      "/vol1/",
      "/vol2/cifs/"
    ],
    "log_path": "/vol0/report_dir",
    "schedule": {
      "name": "weekly"
    },
    "scope": {
      "max_file_size": 1073741824,
      "exclude_paths": [
        "/vol/cold-files/"
      ],
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "exclude_extensions": [
        "mp3"
      ],
      "scan_without_extension": true
    }
```

```
  }
]
}
```

### Creating a Vscan On-Demand policy where a number of optional fields are not specified

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-demand-policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"log_path\":
\"/report\", \"name\": \"on-demand-policy\", \"scan_paths\": [
\"/vol1/cifs/\" ], \"scope\": { \"include_extensions\": [ \"mp*\" ],
\"scan_without_extension\": true }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "name": "on-demand-policy",
    "scan_paths": [
      "vol1/cifs/"
    ],
    "log_path": "/report",
    "scope": {
      "max_file_size": 10737418240,
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "scan_without_extension": true
    }
  }
]
}
```

### Updating a Vscan On-Demand policy

The policy being modified is identified by the UUID of the SVM and the policy name.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-demand-policies/on-demand-policy" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"schedule\":
{ \"name\": \"weekly\" }, \"scope\": { \"exclude_extensions\": [ \"mp3\"
], \"exclude_paths\": [ \"/vol/\" ], \"include_extensions\": [ \"vmdk\",
\"mp3\" ], \"scan_without_extension\": true }}"
```

### Deleting a Vscan On-Demand policy

The policy to be deleted is identified by the UUID of the SVM and the policy name.

```
# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-demand-policies/on-demand-policy" -H "accept:
application/hal+json"
```

### Retrieve a Vscan On-Demand policy

GET `/protocols/vscan/{svm.uuid}/on-demand-policies`

**Introduced In:** 9.6

Retrieves the Vscan On-Demand policy.

#### Related ONTAP commands

- `vserver vscan on-demand-task show`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-demand-policies

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| log_path | string | query | False | Filter by log_path |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.10 |
| schedule.name | string | query | False | Filter by schedule.name |
| schedule.uuid | string | query | False | Filter by schedule.uuid |
| name | string | query | False | Filter by name<br>• maxLength: 256<br>• minLength: 1 |
| scan_paths | string | query | False | Filter by scan_paths |
| scope.include_extensions | string | query | False | Filter by scope.include_extensions |
| scope.max_file_size | integer | query | False | Filter by scope.max_file_size<br>• Max value: 1099511627776<br>• Min value: 1024 |
| scope.exclude_extensions | string | query | False | Filter by scope.exclude_extensions |
| scope.scan_without_extension | boolean | query | False | Filter by scope.scan_without_extension |
| scope.exclude_paths | string | query | False | Filter by scope.exclude_paths<br>• maxLength: 255<br>• minLength: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[vscan_on_demand] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "max_file_size": 10737418240
      },
      "svm": {
        "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

schedule

Schedule of the task.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |

| Name | Type | Description |
| --- | --- | --- |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a Vscan On-Demand policy

POST `/protocols/vscan/{svm.uuid}/on-demand-policies`

**Introduced In:** 9.6

Creates a Vscan On-Demand policy. Created only on a data SVM.

Important notes:

- Only one policy can be scheduled at a time on an SVM. Use schedule name or schedule uuid to schedule an On-Demand policy.
- Scanning must be enabled on the SVM before the policy is scheduled to run.
- The exclude_extensions setting overrides the include_extensions setting. Set scan_without_extension to true to scan files without extensions.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the Vscan On-Demand policy.
- `name` - Name of the Vscan On-Demand policy. Maximum length is 256 characters.
- `log_path` - Path from the Vserver root where the On-Demand policy report is created.
- `scan_paths` - List of paths that need to be scanned.

#### Recommended optional properties

- `schedule` - Scan schedule. It is recommended to set the schedule property, as it dictates when to scan for viruses.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `include_extensions` - *
- `max_file_size` - *10737418240*
- `scan_without_extension` - *true*

#### Related ONTAP commands

- `vserver vscan on-demand-task create`
- `vserver vscan on-demand-task schedule`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-demand-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "log_path": "/vol0/report_dir",
  "name": "task-1",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": 10737418240
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[vscan_on_demand] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "max_file_size": 10737418240
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027101 | The file size must be in the range 1KB to 1TB |
| 10027107 | The include extensions list cannot be empty. Specify at least one extension for inclusion. |
| 10027164 | An On-Demand policy cannot be scheduled, as the Vscan is disabled. Enable the Vscan and retry the operation. |
| 10027167 | The specified schedule does not exist. Create the schedule or create a policy without specifying the schedule. |
| 10027168 | The specified scan path does not exist. The scan path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use “/” not “\”) |
| 10027169 | The specified scan path is not supported for scanning. |
| 10027173 | The new On-Demand policy cannot be created because the SVM has reached the maximum limit of 10 On-Demand policies. Delete an existing policy in order to create a new policy. |
| 10027174 | The specified exclude path is invalid. The path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\") |
| 10027175 | An On-Demand policy cannot be scheduled as the Vserver is not in an operational state. |
| 10027176 | The log-path specified does not exist. The log path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use “/” not “\”). |
| 10027177 | The log path specified is not supported. |

| 10027178 | The specified task-name is invalid. Valid characters are 0-9, A-Z, a-z, "_", ".", and "-"
| 10027253 | The number of paths specified exceeds the configured maximum number of paths. You cannot specify more than the maximum number of configured paths. |
| 10027254 | The number of extensions specified exceeds the configured maximum number of extensions. You cannot specify more than the maximum number of configured extensions. |
| 10027255 | Another policy is already scheduled. Only one policy per SVM is allowed to be scheduled at any one time. Create a policy without specifying a schedule. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

schedule

Schedule of the task.

| Name | Type | Description |
|---|---|---|
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a Vscan On-Demand configuration

DELETE `/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`

**Introduced In:** 9.6

Deletes the Vscan On-Demand configuration.

#### Related ONTAP commands

- `vserver vscan on-demand-task delete`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-demand-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the Vscan On-Demand configuration for an SVM

GET `/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`

**Introduced In:** 9.6

Retrieves the Vscan On-Demand configuration of an SVM.

#### Related ONTAP commands

- `vserver vscan on-demand-task show`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-demand-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| name | string | On-Demand task name |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "log_path": "/vol0/report_dir",
  "name": "task-1",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": 10737418240
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

schedule

Schedule of the task.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the Vscan On-Demand policy configuration for an SVM

PATCH `/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`

**Introduced In:** 9.6

Updates the Vscan On-Demand policy configuration of an SVM. Use schedule name or schedule UUID to schedule an On-Demand scan.

#### Related ONTAP commands

- `vserver vscan on-demand-task modify`
- `vserver vscan on-demand-task schedule`
- `vserver vscan on-demand-task unschedule`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/on-demand-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| log_path | string | The path from the Vserver root where the task report is created. |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "log_path": "/vol0/report_dir",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": 10737418240
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 10027101 | The file size must be in the range 1KB to 1TB |
| 10027107 | The include extensions list cannot be empty. Specify at least one extension for inclusion. |
| 10027164 | An On-Demand policy cannot be scheduled, as the Vscan is disabled. Enable the Vscan and retry the operation. |
| 10027167 | The specified schedule does not exist. Create the schedule or create a policy without specifying the schedule. |
| 10027168 | The specified scan path does not exist. The scan path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use “/” not “\”) |
| 10027169 | The specified scan path is not supported for scanning. |
| 10027174 | The specified exclude path is invalid. The path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\") |
| 10027175 | An On-Demand policy cannot be scheduled as the SVM is not in an operational state. |
| 10027176 | The log-path specified does not exist. The log path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use “/” not “\”) |
| 10027177 | The log path specified is not supported. |
| 10027253 | The number of paths specified exceeds the configured maximum number of paths. You cannot specify more than the maximum number of configured paths. |
| 10027254 | The number of extensions specified exceeds the configured maximum number of extensions. You cannot specify more than the maximum number of configured extensions. |
| 10027255 | Another policy is already scheduled. Only one policy per SVM is allowed to be scheduled at any one time. Update a policy without specifying a schedule. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

schedule

Schedule of the task.

| Name | Type | Description |
|---|---|---|
| name | string | Job schedule name |
| uuid | string | Job schedule UUID |

scope

| Name | Type | Description |
|---|---|---|
| exclude_extensions | array[string] | List of file extensions for which scanning is not performed. |
| exclude_paths | array[string] | List of file paths for which scanning must not be performed. |
| include_extensions | array[string] | List of file extensions to be scanned. |
| max_file_size | integer | Maximum file size, in bytes, allowed for scanning. |
| scan_without_extension | boolean | Specifies whether or not files without any extension can be scanned. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

| Name | Type | Description |
| --- | --- | --- |
| log_path | string | The path from the Vserver root where the task report is created. |
| scan_paths | array[string] | List of paths that need to be scanned. |
| schedule | schedule | Schedule of the task. |
| scope | scope | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage Vscan scanner-pool configuration

## Manage Vscan scanner-pool configuration

### Overview

A scanner-pool defines the Vscan servers and privileged users that can connect to SVMs and a scanner policy or role determines whether a scanner-pool is active. You can configure a scanner-pool to be used on the local cluster or any other cluster in an MCC/DR setup.

### Examples

## Retrieving all fields for all scanner-pools of an SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/<svm-uuid>/scanner-
pools?fields=*&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "1ccbefd6-6132-11ef-92d9-005056bbd354",
      "name": "vs1"
    },
    "name": "scanner-1",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "primary",
    "cluster": {
      "uuid": "3af387d8-6131-11ef-92d9-005056bbd354",
      "name": "clus1"
    }
  },
  {
    "svm": {
      "uuid": "1ccbefd6-6132-11ef-92d9-005056bbd354",
```

```
      "name": "vs1"
    },
    "name": "scanner-2",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "secondary",
    "cluster": {
      "uuid": "3af387d8-6131-11ef-92d9-005056bbd354",
      "name": "clus1"
    }
  }
],
"num_records": 2
}
```


**Retrieving all scanner-pools with *role* set as *secondary***

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/<svm-uuid>/scanner-
pools?role=secondary&fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06",
      "name": "vs1"
    },
    "name": "scanner-2",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "secondary",
    "cluster": {
      "uuid": "0933f9b5-f226-11e8-9601-0050568ecc06",
      "name": "Cluster3"
    }
  }
],
"num_records": 1
}
```

**Retrieving the specified scanner-pool associated with an SVM**

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-
9601-0050568ecc06/scanner-pools/scanner-1?fields=*" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06",
  "name": "vs1"
},
"name": "scanner-1",
"servers": [
  "1.1.1.1",
  "10.72.204.27"
],
"privileged_users": [
  "cifs\\u1",
  "cifs\\u2"
],
"role": "primary",
"cluster": {
  "uuid": "0933f9b5-f226-11e8-9601-0050568ecc06",
  "name": "Cluster3"
}
}
```

**Creating a scanner-pool for an SVM with all fields specified**

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/b103be27-17b8-11e9-
b451-0050568ecd85/scanner-pools?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"cluster\": {
\"name\": \"Cluster1\", \"uuid\": \"ab746d77-17b7-11e9-b450-0050568ecd85\"
}, \"name\": \"test-scanner\", \"privileged_users\": [ \"cifs\\\\u1\",
\"cifs\\\\u2\" ], \"role\": \"primary\", \"servers\": [ \"1.1.1.1\",
\"10.72.204.27\" ]}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "test-scanner",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "primary",
    "cluster": {
      "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85",
      "name": "Cluster1"
    }
  }
]
}
```

**Creating a scanner-pool for an SVM with an unspecified role and cluster**

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/b103be27-17b8-11e9-
b451-0050568ecd85/scanner-pools" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"name\": \"test-scanner-1\",
\"privileged_users\": [ \"cifs\\\\u1\", \"cifs\\\\u2\" ], \"servers\": [
\"1.1.1.1\", \"10.72.204.27\" ]}"

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "test-scanner-1",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ]
  }
]
}
```

**Updating a scanner-pool for an SVM with all of the fields specified**

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-
9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"cluster\": {
\"name\": \"Cluster3\", \"uuid\": \"0933f9b5-f226-11e8-9601-0050568ecc06\"
}, \"privileged_users\": [ \"cifs\\\\u1\", \"cifs\\\\u2\" ], \"role\":
\"secondary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ]}"
```

**Updating the "role" of a scanner-pool for an SVM**

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-
9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"cluster\": {
\"name\": \"Cluster3\", \"uuid\": \"0933f9b5-f226-11e8-9601-0050568ecc06\"
}, \"role\": \"primary\"}"
```

#### Deleting a scanner-pool for a specified SVM

```
# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-
9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept:
application/json"
```

#### Retrieve the Vscan scanner-pool configuration for an SVM

GET `/protocols/vscan/{svm.uuid}/scanner-pools`

**Introduced In:** 9.6

Retrieves the Vscan scanner-pool configuration of an SVM.

##### Related ONTAP commands

- `vserver vscan scanner-pool show`
- `vserver vscan scanner-pool privileged-users show`
- `vserver vscan scanner-pool servers show`

##### Learn more

- DOC /protocols/vscan/{svm.uuid}/scanner-pools

##### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster.name | string | query | False | Filter by cluster.name |
| cluster.uuid | string | query | False | Filter by cluster.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| privileged_users | string | query | False | Filter by privileged_users |
| servers | string | query | False | Filter by servers |
| role | string | query | False | Filter by role |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.10 |
| name | string | query | False | Filter by name<br>• maxLength: 256<br>• minLength: 1 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[vscan_scanner_pool] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "string",
      "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create the Vscan scanner-pool configuration for an SVM**

POST `/protocols/vscan/{svm.uuid}/scanner-pools`

**Introduced In:** 9.6

Creates a Vscan scanner-pool configuration for a specified SVM. You can create a scanner-pool with all fields specified or only mandatory fields specified.

Important notes:

- A scanner-pool must have servers and privileged users specified.
- If the role or cluster is not specified, the scanner-pool is created on the local cluster with the role set as primary.
  *`Only one of the fields cluster-uuid or cluster-name is required.

**Required properties**

- `svm.uuid` or `svm.name` - Existing SVM in which to create the Vscan configuration.
- `name` - Scanner-pool name.
- `privileged_users` - List of privileged users.

- `servers` - List of server IP addresses or FQDNs.

#### Recommended optional properties

- `role` - Setting a role for a scanner-pool is recommended.
- `cluster` - Passing the cluster name or UUID (or both) in a multi-cluster environment is recommended.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `role` - *primary*
- `cluster.name` - Local cluster name.
- `cluster.uuid` - Local cluster UUID.

#### Related ONTAP commands

- `vserver vscan scanner-pool create`
- `vserver vscan scanner-pool apply-policy`
- `vserver vscan scanner-pool privileged-users add`
- `vserver vscan scanner-pool servers add`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "string",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[vscan_scanner_pool] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "string",
      "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 10027076 + | The new scanner-pool cannot be created because the SVM has reached the maximum limit of 20 scanner-pools. Delete an existing scanner-pool in order to create a new scanner-pool |
| 10027086 + | The specified list of servers contains one or more entries that cannot be resolved |
| 10027258 + | The specified cluster_name does not exist |
| 10027256 + | The specified cluster_uuid does not exist |
| 10027257 + | The specified cluster_name and cluster_uuid are valid but belong to different clusters |
| 10027248 + | Scanner-pool created successfully but failed to activate |
| 10027107 + | The list of privileged users or list of servers specified is empty |
| 10027108 + | The list of privileged users specified contains an invalid entry |
| 10027063 + | Attempting to modify a scanner-pool on an SVM owned by the cluster with a data SVM |
| 10027086 + | DNS resolution failed for one or more hostnames |
| 10027119 + | The privileged user contains characters that are not allowed |
| 10027261 + | One of the specified IPv4 addresses in the "-hostnames" is either a multicast, loopback, broadcast, or '0.0.0.0' address. |
| 10027262 + | One of the specified IPv6 addresses in the "-hostnames" is either a link-local, multicast, v4-compatible, v4-mapped, loopback or '::/0' address. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster_reference

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |

| Name | Type | Description |
| --- | --- | --- |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the Vscan scanner-pool configuration

DELETE `/protocols/vscan/{svm.uuid}/scanner-pools/{name}`

**Introduced In:** 9.6

Deletes a Vscan scanner-pool configuration.

Important notes:

- The Vscan scanner-pool DELETE endpoint deletes all of the Vscan scanner-pools for a specified SVM.
- If a Vscan is enabled, it requires at least one scanner-pool to be in the active state. Therefore, disable Vscan on the specified SVM so all the scanner-pools configured on that SVM can be deleted.

#### Related ONTAP commands

- `vserver vscan scanner-pool delete`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027070 + | Attempting to delete a scanner-pool but it is the only active scanner-pool for a Vscan enabled on the SVM |
| 10027064 + | Attempting to delete a scanner-pool with a data SVM which was created by the SVM owned by the cluster |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the Vscan scanner-pool configuration for an SVM

GET `/protocols/vscan/{svm.uuid}/scanner-pools/{name}`

**Introduced In:** 9.6

Retrieves the configuration of a specified scanner-pool of an SVM.

#### Related ONTAP commands

- `vserver vscan scanner-pool show`
- `vserver vscan scanner-pool privileged-users show`
- `vserver vscan scanner-pool servers show`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| name | string | Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".". |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |

| Name | Type | Description |
| --- | --- | --- |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "string",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster_reference

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the Vscan scanner-pool configuration for an SVM

PATCH `/protocols/vscan/{svm.uuid}/scanner-pools/{name}`

**Introduced In:** 9.6

Updates the Vscan scanner-pool configuration of an SVM.

Important notes:

- Along with servers and privileged-users, the role of a scanner-pool can also be updated with the cluster on which a scanner-pool is allowed.
- If role is specified and cluster isn't, then role is applied to the local cluster.

#### Related ONTAP commands

- `vserver vscan scanner-pool modify`
- `vserver vscan scanner-pool apply-policy`
- `vserver vscan scanner-pool privileged-users add`
- `vserver vscan scanner-pool privileged-users remove`
- `vserver vscan scanner-pool servers remove`
- `vserver vscan scanner-pool servers add`

#### Learn more

- DOC /protocols/vscan/{svm.uuid}/scanner-pools

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "string",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 10027258 + | The specified cluster_name does not exist |
| 10027256 + | The specified cluster_uuid does not exist |
| 10027257 + | The specified cluster_name and cluster_uuid are valid but belong to different clusters |
| 10027248 + | Scanner-pool updated successfully but failed to apply the specified role |

| Error Code | Description |
| --- | --- |
| 10027107 + | The list of privileged users or list of servers specified is empty |
| 10027108 + | The list of privileged users specified contains an invalid entry |
| 10027063 + | Attempting to modify a scanner-pool on an SVM owned by the cluster with a data SVM |
| 10027119 + | The privileged user contains characters that are not allowed |
| 10027261 + | One of the specified IPv4 addresses in the "-hostnames" is either a multicast, loopback, broadcast, or '0.0.0.0' address. |
| 10027262 + | One of the specified IPv6 addresses in the "-hostnames" is either a link-local, multicast, v4-compatible, v4-mapped, loopback or '::/0' address. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster_reference

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vscan_scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

| Name | Type | Description |
|---|---|---|
| cluster | cluster_reference | |

| Name | Type | Description |
| --- | --- | --- |
| privileged_users | array[string] | Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.<br>• example: ["cifs\u1", "cifs\u2"]<br>• Introduced in: 9.6 |
| role | string | Specifies the role of the scanner pool. The possible values are:<br>• primary - Always active.<br>• secondary - Active only when none of the primary external virus-scanning servers are connected.<br>• idle - Always inactive. |
| servers | array[string] | Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.<br>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]<br>• Introduced in: 9.6 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# NDMP

## ONTAP REST API NDMP endpoints

NDMP (Network Data Management Protocol) provides an open standard for network-based backup of network-attached storage (NAS). NDMP minimizes the coding required for different applications by providing standard commands for backing up and restoring file servers. NDMP increases the speed and efficiency of NAS data protection, as data can bypass backup servers and be written directly to tape storage. Starting from ONTAP 8.2, the software supports SVM aware NDMP backups. This helps in optimizing NDMP backup performance by choosing efficient data transfer paths and is fully compatible with integrated non-disruptive operations and volume mobility capabilities of ONTAP software. In order to support this feature, a CAB extension must be implemented by backup vendors.
These APIs allow you to manage the following endpoints:

- ndmp
- ndmp nodes
- ndmp svms
- ndmp sessions

### APIs

#### NDMP

These APIs are used to manage NDMP mode, SVM-scope or node-scope.

#### NDMP nodes

These APIs are used to manage node configurations.

#### NDMP SVMs

These APIs are used to manage NDMP configurations of SVMs.

#### NDMP sessions

These APIs are used to manage diagnostic information on NDMP sessions belonging to a specific SVM in the case of SVM-scope or to a specific node in the case of node-scope.

## Manage NDMP mode

### Manage NDMP mode

You can use this API to manage NDMP mode: SVM-scope or node-scope.

### Examples

Updates NDMP mode to SVM:

```
PATCH "/api/protocols/ndmp" '{"mode":"svm"}'
```

Updates NDMP mode to node:

```
PATCH "/api/protocols/ndmp" '{"mode":"node"}'
```

### Retrieve the current NDMP mode

GET `/protocols/ndmp`

**Introduced In:** 9.7

Retrieves the current NDMP mode.

#### Related ONTAP commands

- `system services ndmp node-scope-mode status`

#### Learn more

- DOC /protocols/ndmp

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| mode | string | query | False | Filter by mode |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| mode | string | Indicates whether NDMP is in node-scoped or SVM-scoped mode. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "mode": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the NDMP mode

PATCH `/protocols/ndmp`

**Introduced In:** 9.7

Updates the NDMP mode.

#### Related ONTAP commands

- `system services ndmp node-scope-mode`

#### Learn more

- DOC /protocols/ndmp

#### Request Body

| Name | Type | Description |
|---|---|---|
| mode | string | Indicates whether NDMP is in node-scoped or SVM-scoped mode. |

### Example request

```
{
  "mode": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| mode | string | Indicates whether NDMP is in node-scoped or SVM-scoped mode. |

### Example response

```
{
  "mode": "string"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 65601575 | NDMP node-scope mode is not supported in FSx for ONTAP. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

cluster_ndmp_properties

| Name | Type | Description |
| --- | --- | --- |
| mode | string | Indicates whether NDMP is in node-scoped or SVM-scoped mode. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage NDMP node configurations

## Manage NDMP node configurations

You can use this API to manage node configurations. The authentication types "plaintext" and "plaintext_sso" are used to show that the password uses clear text. Also, they contain no differences for NDMP node scope.

### Examples

Updates "enabled" and "authentication_types" fields:

```
 PATCH "/api/protocols/ndmp/nodes/13bb2092-458b-11e9-9c06-0050568ea64e"
'{"enabled":"false","authentication_types":["plaintext"]}'
```

Updates the "user" field:

```
 PATCH "/api/protocols/ndmp/nodes/13bb2092-458b-11e9-9c06-0050568ea64e"
'{"user":"user22"}'
```

## Retrieve the NDMP node configurations for all nodes

GET `/protocols/ndmp/nodes`

**Introduced In:** 9.7

Retrieves NDMP node configurations for all of the nodes.

#### Related ONTAP commands

- `system services ndmp show`

#### Learn more

- DOC /protocols/ndmp/nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| authentication_types | string | query | False | Filter by authentication_types |
| enabled | boolean | query | False | Filter by enabled |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| user | string | query | False | Filter by user |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of Records |
| records | array[ndmp_node] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_types": [
        "plaintext",
        "challenge"
      ],
      "enabled": 1,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "user": "ndmp_user"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 68812801 | Node-scoped operations are not allowed in an SVM-scope. |
| 68812804 | Failed to get the node name from the specified node UUID. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Storage system node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

ndmp_node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled? |
| node | node | Storage system node |
| user | string | NDMP user ID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the NDMP node configuration for a node

GET `/protocols/ndmp/nodes/{node.uuid}`

**Introduced In:** 9.7

Retrieves an NDMP node configuration for a specific node.

#### Related ONTAP commands

- `system services ndmp show`

#### Learn more

- DOC /protocols/ndmp/nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled? |
| node | node | Storage system node |
| user | string | NDMP user ID |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "user": "ndmp_user"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 68812801 | Node-scoped operations are not allowed in an SVM-scope. |
| 68812802 | The UUID is not valid. |
| 68812804 | Failed to get the node name from the specified node UUID. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Storage system node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the NDMP node configuration for a node

PATCH `/protocols/ndmp/nodes/{node.uuid}`

**Introduced In:** 9.7

Updates the NDMP node configuration for a specific node.

#### Related ONTAP commands

- `system services ndmp modify`

#### Learn more

- DOC /protocols/ndmp/nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled? |
| node | node | Storage system node |
| password | string | NDMP password. This can only be set and cannot be read back. |
| user | string | NDMP user ID |

### Example request

```
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "password": "string",
  "user": "ndmp_user"
}
```

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled? |
| node | node | Storage system node |
| password | string | NDMP password. This can only be set and cannot be read back. |
| user | string | NDMP user ID |

#### Example response

```
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "password": "string",
  "user": "ndmp_user"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 68812800 | The user is required to enable NDMP on a node. |
| 68812801 | Node-scoped operations are not allowed in an SVM-scope. |
| 68812802 | The UUID is not valid. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

Storage system node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

ndmp_node

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled? |
| node | node | Storage system node |
| password | string | NDMP password. This can only be set and cannot be read back. |
| user | string | NDMP user ID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NDMP session configurations

### Manage NDMP session configurations

You can use this API to manage diagnostic information on NDMP sessions belonging to a specific SVM in the case of SVM-scope or to a specific node in the case of node-scope.

### Examples

Retrieves NDMP session details under node-scope:

```
GET "/api/protocols/ndmp/sessions/9b372ce7-3a4b-11e9-a7f8-
0050568e3d73/2000"
```

Retrieves NDMP session details under SVM-scope:

```
GET "/api/protocols/ndmp/sessions/13bb2092-458b-11e9-9c06-
0050568ea604/2000:4000"
```

Deletes NDMP session details under node-scope:

```
DELETE "/api/protocols/ndmp/sessions/9b372ce7-3a4b-11e9-a7f8-
0050568e3d73/2000"
```

Deletes NDMP session details under SVM-scope:

```
DELETE "/api/protocols/ndmp/sessions/13bb2092-458b-11e9-9c06-
0050568ea604/2000:4000"
```

### Retrieve NDMP sessions

GET `/protocols/ndmp/sessions`

**Introduced In:** 9.7

Retrieves a collection of NDMP sessions. In the case of SVM-scope, if this API is executed on a data IP, it displays the list of NDMP sessions under the specified SVM; otherwise it displays the list of NDMP sessions for

all the SVMs under the cluster. In the case of node-scope, it displays the list of NDMP sessions for all nodes.

#### Related ONTAP commands

- `vserver services ndmp probe`
- `system services ndmp probe`

#### Learn more

- DOC /protocols/ndmp/sessions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| client_port | integer | query | False | Filter by client_port |
| uuid | string | query | False | Filter by uuid<br>• Introduced in: 9.9 |
| data.connection.address | string | query | False | Filter by data.connection.address |
| data.connection.port | integer | query | False | Filter by data.connection.port |
| data.connection.type | string | query | False | Filter by data.connection.type |
| data.bytes_processed | integer | query | False | Filter by data.bytes_processed |
| data.operation | string | query | False | Filter by data.operation |
| data.state | string | query | False | Filter by data.state |
| data.reason | string | query | False | Filter by data.reason |
| data_path | string | query | False | Filter by data_path |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| tape_mode | string | query | False | Filter by tape_mode |
| client_address | string | query | False | Filter by client_address |
| mover.reason | string | query | False | Filter by mover.reason |
| mover.state | string | query | False | Filter by mover.state |
| mover.mode | string | query | False | Filter by mover.mode |
| mover.connection.address | string | query | False | Filter by mover.connection.address |
| mover.connection.port | integer | query | False | Filter by mover.connection.port |
| mover.connection.type | string | query | False | Filter by mover.connection.type |
| mover.bytes_moved | integer | query | False | Filter by mover.bytes_moved |
| id | string | query | False | Filter by id |
| backup_engine | string | query | False | Filter by backup_engine |
| tape_device | string | query | False | Filter by tape_device |
| scsi.lun_id | integer | query | False | Filter by scsi.lun_id |
| scsi.host_adapter | integer | query | False | Filter by scsi.host_adapter |
| scsi.device_id | string | query | False | Filter by scsi.device_id |
| scsi.target_id | integer | query | False | Filter by scsi.target_id |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| source_address | string | query | False | Filter by source_address |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of Records. |
| records | array[ndmp_session] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "backup_engine": "string",
      "client_address": "string",
      "data": {
        "bytes_processed": 5000,
        "connection": {
          "address": "string",
          "port": 18600,
          "type": "local"
        },
        "operation": "backup",
        "reason": "aborted",
        "state": "halted"
      },
      "data_path": "/vserver1/vol1",
      "id": "string",
      "mover": {
        "bytes_moved": 645120,
        "connection": {
          "address": "string",
          "port": 18600,
          "type": "local"
        },
        "mode": "read",
        "reason": "end_of_media",
        "state": "connected"
      },
      "node": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scsi": {
        "device_id": "string"
      },
      "source_address": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tape_device": "nrst0a",
      "tape_mode": "write",
      "uuid": "string"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 68812805 | Failed to obtain the NDMP mode of the operation. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ndmp_connect

Indicates the NDMP connection attributes.

| Name | Type | Description |
|---|---|---|
| address | string | Indicates the NDMP data connection address. |
| port | integer | Indicates the NDMP data connection port. |
| type | string | The different NDMP connection types. |

ndmp_data

Information about the NDMP data server.

| Name | Type | Description |
|---|---|---|
| bytes_processed | integer | Indicates the NDMP data bytes processed. |
| connection | ndmp_connect | Indicates the NDMP connection attributes. |
| operation | string | Indicates the NDMP data server operation. |

| Name | Type | Description |
| --- | --- | --- |
| reason | string | If the NDMP mover or data server is in the paused or halted state, then this property provides additional information for the state. |
| state | string | The various states of the NDMP state machine during a backup or restore workflow. |

ndmp_mover

Information about the NDMP mover.

| Name | Type | Description |
| --- | --- | --- |
| bytes_moved | integer | Indicates the NDMP mover bytes moved. |
| connection | ndmp_connect | Indicates the NDMP connection attributes. |
| mode | string | The modes of a NDMP mover instance. |
| reason | string | If the NDMP mover or data server is in the paused or halted state, then this property provides additional information for the state. |
| state | string | The various states of the NDMP state machine during a backup or restore workflow. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

ndmp_scsi

Information about the NDMP SCSI server.

| Name | Type | Description |
| --- | --- | --- |
| device_id | string | Indicates the NDMP SCSI device ID. |
| host_adapter | integer | Indicates the NDMP SCSI host adapter. |
| lun_id | integer | Indicates the NDMP SCSI LUN ID. |
| target_id | integer | Indicates the NDMP SCSI target ID. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ndmp_session

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| backup_engine | string | Indicates the NDMP backup engine. |
| client_address | string | Indicates the NDMP client address. |
| client_port | integer | Indicates the NDMP client port. |
| data | ndmp_data | Information about the NDMP data server. |
| data_path | string | Indicates the NDMP backup or restore path. |

| Name | Type | Description |
| --- | --- | --- |
| id | string | NDMP session identifier. |
| mover | ndmp_mover | Information about the NDMP mover. |
| node | node | |
| scsi | ndmp_scsi | Information about the NDMP SCSI server. |
| source_address | string | Indicates the NDMP local address on which connection was established. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tape_device | string | Indicates the NDMP tape device. |
| tape_mode | string | The modes of a NDMP mover instance. |
| uuid | string | The NDMP node or SVM UUID based on whether NDMP is operating in node-scope or SVM-scope mode. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete an NDMP session

DELETE `/protocols/ndmp/sessions/{owner.uuid}/{session.id}`

**Introduced In:** 9.7

Deletes a specific NDMP session.

#### Related ONTAP commands

- `vserver services ndmp kill`
- `system services ndmp kill`

#### Learn more

- DOC /protocols/ndmp/sessions

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| owner.uuid | string | path | True | NDMP node UUID |
| session.id | string | path | True | Session identifier |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 68812802 | The UUID is not valid. |
| 68812803 | Failed to get the SVM name from the specified SVM UUID. |

| Error code | Description |
| --- | --- |
| 68812804 | Failed to get the node name from the specified node UUID. |
| 68812805 | Failed to obtain the NDMP mode of operation. |
| 68812806 | UUID and Session ID are required. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve NDMP session details

GET `/protocols/ndmp/sessions/{owner.uuid}/{session.id}`

**Introduced In:** 9.7

Retrieves the details of a specific NDMP session.

#### Related ONTAP commands

- `vserver services ndmp probe`
- `system services ndmp probe`

#### Learn more

- DOC /protocols/ndmp/sessions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | The NDMP node or SVM UUID based on whether NDMP is operating in node-scope or SVM-scope mode. |
| session.id | string | path | True | NDMP session identifier |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| backup_engine | string | Indicates the NDMP backup engine. |
| client_address | string | Indicates the NDMP client address. |
| client_port | integer | Indicates the NDMP client port. |
| data | ndmp_data | Information about the NDMP data server. |
| data_path | string | Indicates the NDMP backup or restore path. |
| id | string | NDMP session identifier. |
| mover | ndmp_mover | Information about the NDMP mover. |
| node | node | |
| scsi | ndmp_scsi | Information about the NDMP SCSI server. |

| Name | Type | Description |
| --- | --- | --- |
| source_address | string | Indicates the NDMP local address on which connection was established. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tape_device | string | Indicates the NDMP tape device. |
| tape_mode | string | The modes of a NDMP mover instance. |
| uuid | string | The NDMP node or SVM UUID based on whether NDMP is operating in node-scope or SVM-scope mode. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "backup_engine": "string",
  "client_address": "string",
  "data": {
    "bytes_processed": 5000,
    "connection": {
      "address": "string",
      "port": 18600,
      "type": "local"
    },
    "operation": "backup",
    "reason": "aborted",
    "state": "halted"
  },
  "data_path": "/vserver1/vol1",
  "id": "string",
  "mover": {
    "bytes_moved": 645120,
    "connection": {
      "address": "string",
      "port": 18600,
      "type": "local"
    },
    "mode": "read",
    "reason": "end_of_media",
    "state": "connected"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scsi": {
    "device_id": "string"
  },
```

```
  "source_address": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tape_device": "nrst0a",
  "tape_mode": "write",
  "uuid": "string"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 68812802 | The UUID is not valid. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ndmp_connect

Indicates the NDMP connection attributes.

| Name | Type | Description |
|---|---|---|
| address | string | Indicates the NDMP data connection address. |
| port | integer | Indicates the NDMP data connection port. |
| type | string | The different NDMP connection types. |

ndmp_data

Information about the NDMP data server.

| Name | Type | Description |
|---|---|---|
| bytes_processed | integer | Indicates the NDMP data bytes processed. |
| connection | ndmp_connect | Indicates the NDMP connection attributes. |
| operation | string | Indicates the NDMP data server operation. |
| reason | string | If the NDMP mover or data server is in the paused or halted state, then this property provides additional information for the state. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The various states of the NDMP state machine during a backup or restore workflow. |

ndmp_mover

Information about the NDMP mover.

| Name | Type | Description |
| --- | --- | --- |
| bytes_moved | integer | Indicates the NDMP mover bytes moved. |
| connection | ndmp_connect | Indicates the NDMP connection attributes. |
| mode | string | The modes of a NDMP mover instance. |
| reason | string | If the NDMP mover or data server is in the paused or halted state, then this property provides additional information for the state. |
| state | string | The various states of the NDMP state machine during a backup or restore workflow. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

ndmp_scsi

Information about the NDMP SCSI server.

| Name | Type | Description |
| --- | --- | --- |
| device_id | string | Indicates the NDMP SCSI device ID. |

| Name | Type | Description |
| --- | --- | --- |
| host_adapter | integer | Indicates the NDMP SCSI host adapter. |
| lun_id | integer | Indicates the NDMP SCSI LUN ID. |
| target_id | integer | Indicates the NDMP SCSI target ID. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage NDMP configuration of SVM

## Manage NDMP configuration of SVM

You can use this API to manage NDMP configurations of SVMs.

### Examples

Updates the "enabled" field:

```
 PATCH "/api/protocols/ndmp/svms/9b372ce7-3a4b-11e9-a7f8-0050568e3d73"
'{"enabled":"false"}'
```

Updates the "authentication_types" field:

```
 PATCH "/api/protocols/ndmp/svms/9b372ce7-3a4b-11e9-a7f8-0050568e3d73"
'{"authentication_types":["challenge"]}'
```

## Retrieve NDMP configurations for all SVMs

GET `/protocols/ndmp/svms`

**Introduced In:** 9.7

Retrieves NDMP configurations for all SVMs.

#### Related ONTAP commands

- `vserver services ndmp show`

#### Learn more

- DOC /protocols/ndmp/svms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| authentication_types | string | query | False | Filter by authentication_types |
| enabled | boolean | query | False | Filter by enabled |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of Records |
| records | array[ndmp_svm] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_types": [
        "plaintext",
        "challenge"
      ],
      "enabled": 1,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 2 | The UUID provided is an invalid value for field "svm.uuid". |
| 262197 | The value provided is invalid for field "fields". |
| 65601536 | The operation is not supported because NDMP SVM-aware mode is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ndmp_svm

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled on the SVM? |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve NDMP configuration for an SVM

GET `/protocols/ndmp/svms/{svm.uuid}`

**Introduced In:** 9.7

Retrieves an NDMP configuration for a specific SVM.

#### Related ONTAP commands

- `vserver services ndmp show`

#### Learn more

- DOC /protocols/ndmp/svms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled on the SVM? |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Error**

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 2 | The UUID provided is an invalid value for field "svm.uuid". |
| 262197 | The value provided is invalid for field "fields". |
| 65601536 | The operation is not supported because NDMP SVM-aware mode is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update NDMP configuration for an SVM

PATCH `/protocols/ndmp/svms/{svm.uuid}`

**Introduced In:** 9.7

Updates the NDMP configuration for a specific SVM.

#### Related ONTAP commands

- `vserver services ndmp modify`

#### Learn more

- DOC /protocols/ndmp/svms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled on the SVM? |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled on the SVM? |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
|---|---|
| 65601536 | The operation is not supported because NDMP SVM-aware mode is disabled. |
| 65601551 | Authentication type "plaintext_sso" cannot be combined with other authentication types. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ndmp_svm

| Name | Type | Description |
|---|---|---|
| authentication_types | array[string] | NDMP authentication types. |
| enabled | boolean | Is the NDMP service enabled on the SVM? |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Generate the NDMP user password

### Generate the NDMP user password

Generates and retrieves the password for a given NDMP user in the SVM context.

### Examples

```
GET "/api/protocols/ndmp/svms/ca8e29e0-e116-11ea-876c-
0050568ea754/passwords/ndmpuser"
```

### Generate and retrieve the password for an NDMP user

GET `/protocols/ndmp/svms/{svm.uuid}/passwords/{user}`

**Introduced In:** 9.9

Generates and retrieves the password for the specified NDMP user.

#### Related ONTAP commands

- `vserver services ndmp generate-password`

#### Learn more

- DOC /protocols/ndmp/svms/{svm.uuid}/passwords/{user}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| user | string | path | True | User name |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| password | string | NDMP Password |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| user | string | NDMP user |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "password": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user": "string"
}
```


#### Error

```
Status: Default
```

ONTAP Error Response codes

| Error code | Description |
| --- | --- |
| 65601536 | The operation is not supported because NDMP SVM-aware mode is disabled. |
| 65601560 | An NDMP-specific password is not supported for authentication type "plaintext_sso". |
| 65601587 | The NDMP protocol is not supported on this platform. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# NVMe

## ONTAP REST API NVMe endpoints

### Overview

The Non-Volatile Memory Express (NVMe) API endpoints and objects provide for configuration, provisioning and management of the NVMe-related objects.
NVMe over Fabrics (NVMe-oF) refers to the extensions and changes to the base NVMe command set to support NVMe commands over a fabric interconnect and from multiple hosts simultaneously. ONTAP implements elements of both NVMe and NVMe-oF. Throughout this documentation, NVMe is generally used to refer to both NVMe and NVMe-oF.

#### Fibre Channel Logins

Fibre Channel logins represent connections, formed by Fibre Channel initiators, that have successfully logged in to ONTAP. This represents the Fibre Channel login on which higher-level protocols such as Fibre Channel Protocol (FCP) and Non-Volatile Memory Express over Fibre Channel (NVMe over FC) rely.

The Fibre Channel logins REST API provides information about active Fibre Channel logins.

#### NVMe Interfaces

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics protocol. The NVMe interfaces are Fibre Channel interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for the host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

The NVMe interfaces REST API provides NVMe-specific information about network interfaces configured to support an NVMe-oF protocol.

### Learn More

- *Fibre Channel Interfaces* found in the *networking* section. Fibre Channel interfaces are the logical endpoints for Fibre Channel network connections to an SVM.

#### NVMe Services

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for a given SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

The Non-Volatile Memory Express (NVMe) service REST API allows you to create, update, delete, and discover NVMe services for SVMs.

#### NVMe Subsystem Controllers

Non-Volatile Memory Express (NVMe) subsystem controllers represent dynamic connections between hosts and a storage solution.

The NVMe subsystem controllers REST API provides information about connected hosts.

#### NVMe Subsystem Maps

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

The NVMe subsystem map REST API allows you to create, delete, and discover NVMe subsystem maps.

#### NVMe Subsystems

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

The NVMe subsystem REST API allows you to create, update, delete, and discover NVMe subsystems. It also allows you to add and remove NVMe hosts that can access the subsystem and associated namespaces.

#### NVMe Namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

The NVMe namespace REST API allows you to create, update, delete, and discover NVMe namespaces.

## View NVMe interfaces

### View NVMe interfaces

#### Overview

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics (NVMe-oF) protocol. The NVMe interfaces are Fibre Channel (FC) interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

The NVMe interfaces REST API provides NVMe-specific information about network interfaces configured to support an NVMe-oF protocol.

NVMe interfaces must be created using the protocol-specific endpoints for FC interfaces. See POST /network/fc/interfaces . After creation, the interfaces are available via this interface.

#### Examples

### Retrieving summary information for all NVMe interfaces

```
# The API:
GET /api/protocols/nvme/interfaces

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/interfaces' -H 'Accept:
application/hal+json'
```

```
# The response:
{
"records": [
  {
    "svm": {
      "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
        }
      }
    },
    "uuid": "74d69872-0d30-11e9-a684-005056bbdb14",
    "name": "nvme1",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/interfaces/74d69872-0d30-11e9-a684-
005056bbdb14"
      }
    }
  },
  {
    "svm": {
      "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
        }
      }
    },
    "uuid": "77ded991-0d30-11e9-a684-005056bbdb14",
    "name": "nvme2",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/interfaces/77ded991-0d30-11e9-a684-
005056bbdb14"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/interfaces"
```

```
  }
}
}
```

#### Retrieving detailed information for a specific NVMe interface

```
# The API:
GET /api/protocols/nvme/interfaces/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/interfaces/77ded991-
0d30-11e9-a684-005056bbdb14' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
    }
  }
},
"uuid": "77ded991-0d30-11e9-a684-005056bbdb14",
"name": "nvme2",
"enabled": true,
"node": {
  "name": "node1",
  "uuid": "cd4d47fd-0d2e-11e9-a684-005056bbdb14",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/cd4d47fd-0d2e-11e9-a684-005056bbdb14"
    }
  }
},
"transport_address": "nn-0x2003005056bbdb14:pn-0x2005005056bbdb14",
"fc_interface": {
  "wwnn": "20:03:00:50:56:bb:db:14",
  "wwpn": "20:05:00:50:56:bb:db:14",
  "port": {
    "name": "1a",
    "uuid": "081ec491-0d2f-11e9-a684-005056bbdb14",
    "node": {
```

```
      "name": "node1"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/081ec491-0d2f-11e9-a684-
005056bbdb14"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/network/fc/interfaces/77ded991-0d30-11e9-a684-
005056bbdb14"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/protocols/nvme/interfaces/77ded991-0d30-11e9-a684-
005056bbdb14"
  }
}
}
```


### Retrieve NVMe interfaces

GET `/protocols/nvme/interfaces`

**Introduced In:** 9.6

Retrieves NVMe interfaces.

#### Related ONTAP commands

- `vserver nvme show-interface`

#### Learn more

- DOC /protocols/nvme/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| transport_address | string | query | False | Filter by transport_address |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | query | False | Filter by uuid |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| interface_type | string | query | False | Filter by interface_type<br><br>• Introduced in: 9.10 |
| ip_interface.location.port.node.name | string | query | False | Filter by ip_interface.location.port.node.name<br><br>• Introduced in: 9.10 |
| ip_interface.location.port.uuid | string | query | False | Filter by ip_interface.location.port.uuid<br><br>• Introduced in: 9.10 |
| ip_interface.location.port.name | string | query | False | Filter by ip_interface.location.port.name<br><br>• Introduced in: 9.10 |
| ip_interface.ip.address | string | query | False | Filter by ip_interface.ip.address<br><br>• Introduced in: 9.10 |
| enabled | boolean | query | False | Filter by enabled |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fc_interface.wwpn | string | query | False | Filter by fc_interface.wwpn |
| fc_interface.wwnn | string | query | False | Filter by fc_interface.wwnn |
| fc_interface.port.name | string | query | False | Filter by fc_interface.port.name |
| fc_interface.port.node.name | string | query | False | Filter by fc_interface.port.node.name |
| fc_interface.port.uuid | string | query | False | Filter by fc_interface.port.uuid |
| name | string | query | False | Filter by name |
| transport_protocols | string | query | False | Filter by transport_protocols<br>• Introduced in: 9.10 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "fc_interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "wwnn": "20:00:00:50:56:b4:13:a9",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "interface_type": "string",
      "ip_interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "location": {
          "port": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      },
      "name": "lif1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "transport_address": "nn-0x200a00a0989062da:pn-
0x200100a0989062da",
      "transport_protocols": [
        "string"
      ],
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
```

```
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

fc_interface

The attributes specific to a Fibre Channel-based NVMe interface.

This is populated when `interface_type` is *fc_interface*.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| wwnn | string | The WWNN (world wide node name) of the Fibre Channel NVMe interface. |
| wwpn | string | The WWPN (world wide port name) of the Fibre Channel NVMe interface. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

ip

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port_reference

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

location

| Name | Type | Description |
| --- | --- | --- |
| port | port_reference | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

ip_interface

The attributes specific to an IP-based NVMe interface.

This is populated when `interface_type` is *ip_interface*.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| ip | ip | |
| location | location | |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_interface

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics (NVMe-oF) protocol. The NVMe interfaces are Fibre Channel interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

NVMe interfaces must be created using the protocol-specific endpoints for Fibre Channel interfaces. See POST /network/fc/interfaces . After creation, the interfaces are available via this interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the NVMe interface. |
| fc_interface | fc_interface | The attributes specific to a Fibre Channel-based NVMe interface.<br><br>This is populated when `interface_type` is *fc_interface*. |
| interface_type | string | The underlying interface type of the NVMe interface. This property identifies which of *fc_interface* and *ip_interface* will be further populated. |
| ip_interface | ip_interface | The attributes specific to an IP-based NVMe interface.<br><br>This is populated when `interface_type` is *ip_interface*. |
| name | string | The name of the NVMe interface. |
| node | node | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport_address | string | The transport address of the NVMe interface. |
| transport_protocols | array[string] | The transport protocols supported by the NVMe interface. |
| uuid | string | The unique identifier of the NVMe interface. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe interface

GET `/protocols/nvme/interfaces/{uuid}`

**Introduced In:** 9.6

Retrieves an NVMe interface.

#### Related ONTAP commands

- `vserver nvme show-interface`

#### Learn more

- DOC /protocols/nvme/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe interface. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the NVMe interface. |
| fc_interface | fc_interface | The attributes specific to a Fibre Channel-based NVMe interface.<br><br>This is populated when `interface_type` is *fc_interface*. |
| interface_type | string | The underlying interface type of the NVMe interface. This property identifies which of *fc_interface* and *ip_interface* will be further populated. |
| ip_interface | ip_interface | The attributes specific to an IP-based NVMe interface.<br><br>This is populated when `interface_type` is *ip_interface*. |
| name | string | The name of the NVMe interface. |
| node | node | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| transport_address | string | The transport address of the NVMe interface. |
| transport_protocols | array[string] | The transport protocols supported by the NVMe interface. |
| uuid | string | The unique identifier of the NVMe interface. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "fc_interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "wwnn": "20:00:00:50:56:b4:13:a9",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "interface_type": "string",
  "ip_interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "location": {
      "port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    },
    "name": "lif1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "transport_address": "nn-0x200a00a0989062da:pn-0x200100a0989062da",
    "transport_protocols": [
      "string"
    ],
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The supplied SVM does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

fc_interface

The attributes specific to a Fibre Channel-based NVMe interface.

This is populated when `interface_type` is *fc_interface*.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| wwnn | string | The WWNN (world wide node name) of the Fibre Channel NVMe interface. |
| wwpn | string | The WWPN (world wide port name) of the Fibre Channel NVMe interface. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

ip

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port_reference

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

location

| Name | Type | Description |
| --- | --- | --- |
| port | port_reference | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

ip_interface

The attributes specific to an IP-based NVMe interface.

This is populated when `interface_type` is *ip_interface*.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| ip | ip | |
| location | location | |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NVMe services

### Manage NVMe services

#### Overview

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

The Non-Volatile Memory Express (NVMe) service REST API allows you to create, update, delete, and discover NVMe services for SVMs.

#### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

### Creating an NVMe service for an SVM

The simplest way to create an NVMe service is to specify only the SVM, either by name or UUID. By default, the new NVMe service is enabled.

In this example, the `return_records` query parameter is used to retrieve the new NVMe service object in the REST response.

```
# The API:
POST /api/protocols/nvme/services

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/nvme/services?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
        }
      }
    },
    "enabled": true,
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-
005056bb7341"
      }
    }
  }
]
}
```

#### Retrieving the NVMe services for all SVMs in the cluster

```
# The API:
GET /api/protocols/nvme/services

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services' -H 'Accept:
application/hal+json'

# The response:
```

```
{
"records": [
  {
    "svm": {
      "uuid": "ab60c350-dc68-11e8-9711-005056bbe408",
      "name": "svm0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/ab60c350-dc68-11e8-9711-005056bbe408"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/services/ab60c350-dc68-11e8-9711-
005056bbe408"
      }
    }
  },
  {
    "svm": {
      "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-
005056bb7341"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/services"
  }
}
}
```

### Retrieving details for a specific NVMe service

The NVMe service is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/nvme/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-dc69-
11e8-b29f-005056bb7341' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
    }
  }
},
"enabled": true,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-
005056bb7341"
  }
}
}
```

### Disabling an NVMe service

Disabling an NVMe service shuts down all active NVMe connections for the SVM and prevents the creation of new NVMe connections.

The NVMe service to update is identified by the UUID of its SVM.

```
# The API:
PATCH /api/protocols/nvme/services/{svm.uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-
dc69-11e8-b29f-005056bb7341' -H 'Accept: application/hal+json' -d '{
"enabled": "false" }'
```

You can retrieve the NVMe service to confirm the change.

```
# The API:
GET /api/protocols/nvme/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-dc69-
11e8-b29f-005056bb7341' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
    }
  }
},
"enabled": false,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-
005056bb7341"
  }
}
}
```

---

#### Deleting an NVMe service

The NVMe service must be disabled before it can be deleted. In addition, all NVMe interfaces, subsystems, and subsystem maps associated with the SVM must first be deleted.

The NVMe service to delete is identified by the UUID of its SVM.

```
# The API:
DELETE /api/protocols/nvme/services/{svm.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-
dc69-11e8-b29f-005056bb7341' -H 'Accept: application/hal+json'
```

### Retrieve NVMe services

GET `/protocols/nvme/services`

**Introduced In:** 9.6

Retrieves NVMe services.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver nvme show`

#### Learn more

- DOC /protocols/nvme/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br>• Introduced in: 9.7 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.7 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.7 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.7 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.7 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.7 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.7 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.7 |
| metric.fc.iops.other | integer | query | False | Filter by metric.fc.iops.other<br><br>• Introduced in: 9.10 |
| metric.fc.iops.total | integer | query | False | Filter by metric.fc.iops.total<br><br>• Introduced in: 9.10 |
| metric.fc.iops.write | integer | query | False | Filter by metric.fc.iops.write<br><br>• Introduced in: 9.10 |
| metric.fc.iops.read | integer | query | False | Filter by metric.fc.iops.read<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.fc.latency.other | integer | query | False | Filter by metric.fc.latency.other<br><br>• Introduced in: 9.10 |
| metric.fc.latency.total | integer | query | False | Filter by metric.fc.latency.total<br><br>• Introduced in: 9.10 |
| metric.fc.latency.write | integer | query | False | Filter by metric.fc.latency.write<br><br>• Introduced in: 9.10 |
| metric.fc.latency.read | integer | query | False | Filter by metric.fc.latency.read<br><br>• Introduced in: 9.10 |
| metric.fc.timestamp | string | query | False | Filter by metric.fc.timestamp<br><br>• Introduced in: 9.10 |
| metric.fc.throughput.write | integer | query | False | Filter by metric.fc.throughput.write<br><br>• Introduced in: 9.10 |
| metric.fc.throughput.read | integer | query | False | Filter by metric.fc.throughput.read<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.fc.throughput.total | integer | query | False | Filter by metric.fc.throughput.total<br><br>• Introduced in: 9.10 |
| metric.fc.status | string | query | False | Filter by metric.fc.status<br><br>• Introduced in: 9.10 |
| metric.fc.duration | string | query | False | Filter by metric.fc.duration<br><br>• Introduced in: 9.10 |
| metric.tcp.iops.other | integer | query | False | Filter by metric.tcp.iops.other<br><br>• Introduced in: 9.10 |
| metric.tcp.iops.total | integer | query | False | Filter by metric.tcp.iops.total<br><br>• Introduced in: 9.10 |
| metric.tcp.iops.write | integer | query | False | Filter by metric.tcp.iops.write<br><br>• Introduced in: 9.10 |
| metric.tcp.iops.read | integer | query | False | Filter by metric.tcp.iops.read<br><br>• Introduced in: 9.10 |
| metric.tcp.latency.other | integer | query | False | Filter by metric.tcp.latency.other<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.tcp.latency.total | integer | query | False | Filter by metric.tcp.latency.total<br><br>• Introduced in: 9.10 |
| metric.tcp.latency.write | integer | query | False | Filter by metric.tcp.latency.write<br><br>• Introduced in: 9.10 |
| metric.tcp.latency.read | integer | query | False | Filter by metric.tcp.latency.read<br><br>• Introduced in: 9.10 |
| metric.tcp.timestamp | string | query | False | Filter by metric.tcp.timestamp<br><br>• Introduced in: 9.10 |
| metric.tcp.throughput.write | integer | query | False | Filter by metric.tcp.throughput.write<br><br>• Introduced in: 9.10 |
| metric.tcp.throughput.read | integer | query | False | Filter by metric.tcp.throughput.read<br><br>• Introduced in: 9.10 |
| metric.tcp.throughput.total | integer | query | False | Filter by metric.tcp.throughput.total<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.tcp.status | string | query | False | Filter by metric.tcp.status<br><br>• Introduced in: 9.10 |
| metric.tcp.duration | string | query | False | Filter by metric.tcp.duration<br><br>• Introduced in: 9.10 |
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.7 |
| statistics.fc.latency_raw.other | integer | query | False | Filter by statistics.fc.latency_raw.other<br><br>• Introduced in: 9.10 |
| statistics.fc.latency_raw.total | integer | query | False | Filter by statistics.fc.latency_raw.total<br><br>• Introduced in: 9.10 |
| statistics.fc.latency_raw.write | integer | query | False | Filter by statistics.fc.latency_raw.write<br><br>• Introduced in: 9.10 |
| statistics.fc.latency_raw.read | integer | query | False | Filter by statistics.fc.latency_raw.read<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.fc.iops_raw.other | integer | query | False | Filter by statistics.fc.iops_raw.other<br><br>• Introduced in: 9.10 |
| statistics.fc.iops_raw.total | integer | query | False | Filter by statistics.fc.iops_raw.total<br><br>• Introduced in: 9.10 |
| statistics.fc.iops_raw.write | integer | query | False | Filter by statistics.fc.iops_raw.write<br><br>• Introduced in: 9.10 |
| statistics.fc.iops_raw.read | integer | query | False | Filter by statistics.fc.iops_raw.read<br><br>• Introduced in: 9.10 |
| statistics.fc.status | string | query | False | Filter by statistics.fc.status<br><br>• Introduced in: 9.10 |
| statistics.fc.throughput_raw.write | integer | query | False | Filter by statistics.fc.throughput_raw.write<br><br>• Introduced in: 9.10 |
| statistics.fc.throughput_raw.read | integer | query | False | Filter by statistics.fc.throughput_raw.read<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.fc.throughput_raw.total | integer | query | False | Filter by statistics.fc.throughput_raw.total<br>• Introduced in: 9.10 |
| statistics.fc.timestamp | string | query | False | Filter by statistics.fc.timestamp<br>• Introduced in: 9.10 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.7 |
| statistics.tcp.latency_raw.other | integer | query | False | Filter by statistics.tcp.latency_raw.other<br>• Introduced in: 9.10 |
| statistics.tcp.latency_raw.total | integer | query | False | Filter by statistics.tcp.latency_raw.total<br>• Introduced in: 9.10 |
| statistics.tcp.latency_raw.write | integer | query | False | Filter by statistics.tcp.latency_raw.write<br>• Introduced in: 9.10 |
| statistics.tcp.latency_raw.read | integer | query | False | Filter by statistics.tcp.latency_raw.read<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.tcp.iops_raw.other | integer | query | False | Filter by statistics.tcp.iops_raw.other<br><br>• Introduced in: 9.10 |
| statistics.tcp.iops_raw.total | integer | query | False | Filter by statistics.tcp.iops_raw.total<br><br>• Introduced in: 9.10 |
| statistics.tcp.iops_raw.write | integer | query | False | Filter by statistics.tcp.iops_raw.write<br><br>• Introduced in: 9.10 |
| statistics.tcp.iops_raw.read | integer | query | False | Filter by statistics.tcp.iops_raw.read<br><br>• Introduced in: 9.10 |
| statistics.tcp.status | string | query | False | Filter by statistics.tcp.status<br><br>• Introduced in: 9.10 |
| statistics.tcp.throughput_raw.write | integer | query | False | Filter by statistics.tcp.throughput_raw.write<br><br>• Introduced in: 9.10 |
| statistics.tcp.throughput_raw.read | integer | query | False | Filter by statistics.tcp.throughput_raw.read<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.tcp.throughput_raw.total | integer | query | False | Filter by statistics.tcp.throughput_raw.total<br><br>• Introduced in: 9.10 |
| statistics.tcp.timestamp | string | query | False | Filter by statistics.tcp.timestamp<br><br>• Introduced in: 9.10 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.7 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.7 |
| svm.name | string | query | False | Filter by svm.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | query | False | Filter by svm.uuid |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "fc": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput": {
            "read": 200,
```

```
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "tcp": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
```

```
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "statistics": {
        "fc": {
          "iops_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "tcp": {
          "iops_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency_raw": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "throughput_raw": {
```

```
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| fc | fc | The NVMe/FC portion of the aggregated metrics. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| tcp | tcp | The NVMe/TCP portion of the aggregated metrics. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated statistics.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | The NVMe/FC portion of the aggregated statistics. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| tcp | tcp | The NVMe/TCP portion of the aggregated statistics. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_service

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an NVMe service

POST `/protocols/nvme/services`

**Introduced In:** 9.6

Creates an NVMe service.

#### Required properties

- `svm.uuid` or `svm.name` - The existing SVM in which to create the NVMe service.

#### Related ONTAP commands

- `vserver nvme create`

#### Learn more

- DOC /protocols/nvme/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[nvme_service] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115127 | The cluster lacks a valid NVMe license. |
| 2621462 | The supplied SVM does not exist. |
| 2621507 | NVMe is not allowed for the specified SVM. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |

| Error Code | Description |
| --- | --- |
| 5374893 | The SVM is stopped. The SVM must be running to create an NVMe service. |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |
| 72089650 | An NVMe service already exists for the specified SVM. |
| 72089900 | An NVMe service cannot be creating in an SVM that is configured for a SAN protocol. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_service

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an NVMe service

DELETE `/protocols/nvme/services/{svm.uuid}`

**Introduced In:** 9.6

Deletes an NVMe service. An NVMe service must be disabled before it can be deleted. In addition, all NVMe interfaces, subsystems, and subsystem maps associated with the SVM must first be deleted.

#### Related ONTAP commands

- `vserver nvme delete`

#### Learn more

- DOC /protocols/nvme/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM whose NVMe service is to be deleted. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The supplied SVM does not exist. |

| Error Code | Description |
| --- | --- |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |
| 72089651 | The supplied SVM does not have an NVMe service. |
| 72089653 | There are subsystems associated with the NVMe service SVM. The subsystems must be removed before deleting the NVMe service. |
| 72089654 | There are NVMe-oF LIFs associated with the NVMe service SVM. The LIFs must be removed before deleting the NVMe service. |
| 72090028 | The NVMe service is enabled. The NVMe service must be disabled before it can be deleted. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe service

GET `/protocols/nvme/services/{svm.uuid}`

**Introduced In:** 9.6

Retrieves an NVMe service.

#### Related ONTAP commands

- `vserver nvme show`

#### Learn more

- DOC /protocols/nvme/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM whose NVMe service is to be retrieved. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "fc": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
```

```
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "tcp": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "statistics": {
    "fc": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
```

```
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "tcp": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
```

```
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The supplied SVM does not exist. |
| 72089651 | The supplied SVM does not have an NVMe service. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | The NVMe/FC portion of the aggregated metrics. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| tcp | tcp | The NVMe/TCP portion of the aggregated metrics. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | The NVMe/FC portion of the aggregated statistics. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| tcp | tcp | The NVMe/TCP portion of the aggregated statistics. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an NVMe service

PATCH `/protocols/nvme/services/{svm.uuid}`

**Introduced In:** 9.6

Updates an NVMe service.

#### Related ONTAP commands

- `vserver nvme modify`

#### Learn more

- DOC /protocols/nvme/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM whose NVMe service is to be updated. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1115127 | The cluster lacks a valid NVMe license. |
| 2621462 | The supplied SVM does not exist. |
| 5374893 | The SVM is stopped. The SVM must be running to create an NVMe service. |
| 72089651 | The supplied SVM does not have an NVMe service. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

tcp

The NVMe/TCP portion of the aggregated statistics.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_service

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Retrieve NVMe protocol historical performance metrics

GET `/protocols/nvme/services/{svm.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for the NVMe protocol service of an SVM.

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| fc.status | string | query | False | Filter by fc.status<br>• Introduced in: 9.10 |
| fc.duration | string | query | False | Filter by fc.duration<br>• Introduced in: 9.10 |
| fc.throughput.write | integer | query | False | Filter by fc.throughput.write<br>• Introduced in: 9.10 |
| fc.throughput.read | integer | query | False | Filter by fc.throughput.read<br>• Introduced in: 9.10 |
| fc.throughput.total | integer | query | False | Filter by fc.throughput.total<br>• Introduced in: 9.10 |
| fc.latency.other | integer | query | False | Filter by fc.latency.other<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| fc.latency.total | integer | query | False | Filter by fc.latency.total<br><br>• Introduced in: 9.10 |
| fc.latency.write | integer | query | False | Filter by fc.latency.write<br><br>• Introduced in: 9.10 |
| fc.latency.read | integer | query | False | Filter by fc.latency.read<br><br>• Introduced in: 9.10 |
| fc.iops.other | integer | query | False | Filter by fc.iops.other<br><br>• Introduced in: 9.10 |
| fc.iops.total | integer | query | False | Filter by fc.iops.total<br><br>• Introduced in: 9.10 |
| fc.iops.write | integer | query | False | Filter by fc.iops.write<br><br>• Introduced in: 9.10 |
| fc.iops.read | integer | query | False | Filter by fc.iops.read<br><br>• Introduced in: 9.10 |
| duration | string | query | False | Filter by duration |
| tcp.status | string | query | False | Filter by tcp.status<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| tcp.duration | string | query | False | Filter by tcp.duration<br>• Introduced in: 9.10 |
| tcp.throughput.write | integer | query | False | Filter by tcp.throughput.write<br>• Introduced in: 9.10 |
| tcp.throughput.read | integer | query | False | Filter by tcp.throughput.read<br>• Introduced in: 9.10 |
| tcp.throughput.total | integer | query | False | Filter by tcp.throughput.total<br>• Introduced in: 9.10 |
| tcp.latency.other | integer | query | False | Filter by tcp.latency.other<br>• Introduced in: 9.10 |
| tcp.latency.total | integer | query | False | Filter by tcp.latency.total<br>• Introduced in: 9.10 |
| tcp.latency.write | integer | query | False | Filter by tcp.latency.write<br>• Introduced in: 9.10 |
| tcp.latency.read | integer | query | False | Filter by tcp.latency.read<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| tcp.iops.other | integer | query | False | Filter by tcp.iops.other<br><br>• Introduced in: 9.10 |
| tcp.iops.total | integer | query | False | Filter by tcp.iops.total<br><br>• Introduced in: 9.10 |
| tcp.iops.write | integer | query | False | Filter by tcp.iops.write<br><br>• Introduced in: 9.10 |
| tcp.iops.read | integer | query | False | Filter by tcp.iops.read<br><br>• Introduced in: 9.10 |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| status | string | query | False | Filter by status |
| svm.uuid | string | path | True | The unique identifier of the SVM. |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "fc": {
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        }
      },
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
```

```
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "svm": {
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tcp": {
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        }
      },
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |

| Error Code | Description |
| --- | --- |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

records

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| fc | fc | The NVMe/FC portion of the aggregated metrics. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| tcp | tcp | The NVMe/TCP portion of the aggregated metrics. |

| Name | Type | Description |
|---|---|---|
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for the NVMe protocol service of an SVM for a specific time**

GET `/protocols/nvme/services/{svm.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical performance metrics for the NVMe protocol service of an SVM for a specific time.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| fc | fc | The NVMe/FC portion of the aggregated metrics. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| tcp | tcp | The NVMe/TCP portion of the aggregated metrics. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "fc": {
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    }
  },
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "svm": {
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tcp": {
    "duration": "PT15S",
    "iops": {
      "read": 200,
```

```
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    }
  },
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

fc

The NVMe/FC portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. |

tcp

The NVMe/TCP portion of the aggregated metrics.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View NVMe subsystem controllers

### View NVMe subsystem controllers

#### Overview

Non-Volatile Memory Express (NVMe) subsystem controllers represent dynamic connections between hosts and a storage solution.

The NVMe subsystem controllers REST API provides information about connected hosts.

#### Examples

### Retrieving the NVMe subsystem controllers for the entire system

```
# The API:
GET /api/protocols/nvme/subsystem-controllers

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-controllers'
-H 'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
```

```
      "uuid": "14875240-2594-11e9-abde-00a098984313",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-
00a098984313"
        }
      }
    },
    "id": "0040h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-
11e9-abde-00a098984313/0040h"
      }
    }
  },
  {
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
      "uuid": "14875240-2594-11e9-abde-00a098984313",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-
00a098984313"
        }
      }
    },
    "id": "0041h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-
11e9-abde-00a098984313/0041h"
      }
    }
  },
  {
```

```
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
      "uuid": "1489d0d5-2594-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_1",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/1489d0d5-2594-11e9-94c4-
00a0989a1c8e"
        }
      }
    },
    "id": "0040h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/1489d0d5-2594-
11e9-94c4-00a0989a1c8e/0040h"
      }
    }
  },
  {
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
      "uuid": "1489d0d5-2594-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_1",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/1489d0d5-2594-11e9-94c4-
00a0989a1c8e"
        }
      }
```

```
    },
    "id": "0041h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/1489d0d5-2594-
11e9-94c4-00a0989a1c8e/0041h"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystem-controllers"
  }
}
}
```

#### Retrieving the NVMe subsystem controllers for a specific subsystem

```
# The API:
GET /api/protocols/nvme/subsystem-controllers

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-
controllers?subsystem.uuid=14875240-2594-11e9-abde-00a098984313' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
      "uuid": "14875240-2594-11e9-abde-00a098984313",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
```

```
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-
00a098984313"
        }
      }
    },
    "id": "0040h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-
11e9-abde-00a098984313/0040h"
      }
    }
  },
  {
    "svm": {
      "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_fcnvme_vserver_0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "subsystem": {
      "uuid": "14875240-2594-11e9-abde-00a098984313",
      "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-
00a098984313"
        }
      }
    },
    "id": "0041h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-
11e9-abde-00a098984313/0041h"
      }
    }
  }
],
"num_records": 2,
"_links": {
```

```
  "self": {
    "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-
abde-00a098984313"
  }
}
}
```

#### Retrieving a specific NVMe subsystem controller

```
# The API:
GET /api/protocols/nvme/subsystem-controllers/{subsystem.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-
controllers/14875240-2594-11e9-abde-00a098984313/0040h' -H 'Accept:
application/hal+json'

# The response:
{
"svm": {
  "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
  "name": "symmcon_fcnvme_vserver_0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
    }
  }
},
"subsystem": {
  "uuid": "14875240-2594-11e9-abde-00a098984313",
  "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-
00a098984313"
    }
  }
},
"id": "0040h",
"interface": {
  "name": "symmcon_lif_fcnvme_symmcon_fcnvme_vserver_0_3a_0",
  "uuid": "fa1c5941-2593-11e9-94c4-00a0989a1c8e",
  "transport_address": "nn-0x200400a0989a1c8d:pn-0x200500a0989a1c8d",
  "_links": {
```

```
    "self": {
      "href": "/api/protocols/nvme/interfaces/fa1c5941-2593-11e9-94c4-
00a0989a1c8e"
    }
  }
},
"node": {
  "name": "ssan-8040-94a",
  "uuid": "ebf66f05-2590-11e9-abde-00a098984313",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/ebf66f05-2590-11e9-abde-00a098984313"
    }
  }
},
"host": {
  "transport_address": "nn-0x20000090fae00806:pn-0x10000090fae00806",
  "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-a3b0-
71ce907b4eef",
  "id": "b8546ca6097349e5b1558dc154fc073b"
},
"io_queue": {
  "count": 4,
  "depth": [
    32,
    32,
    32,
    32
  ]
},
"admin_queue": {
  "depth": 32
},
"dh_hmac_chap": {
  "mode": "none"
},
"keep_alive_timeout": 4000,
"digest.header": true,
"digest.data": false,
"tls": {
  "key_type": "configured",
  "psk_identity": "NVMe1R01 nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-
4020-a3b0-71ce907b4eef nqn.1992-
08.com.netapp:sn.ca3cae02070811ef9a53005056bb9001:subsystem.ss1
c9X3RurQxGiGa76Tpk2tirifrUhHmVp035MOrtHXnAU=",
  "cipher": "tls_aes_128_gcm_sha256"
```

```
},
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0040h"
  }
}
}
```


### Retrieve NVMe subsystem controllers

GET `/protocols/nvme/subsystem-controllers`

**Introduced In:** 9.6

Retrieves NVMe subsystem controllers.

#### Related ONTAP commands

- `vserver nvme subsystem controller show`

#### Learn more

- DOC /protocols/nvme/subsystem-controllers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| subsystem.uuid | string | query | False | Filter by subsystem.uuid |
| subsystem.name | string | query | False | Filter by subsystem.name<br>• maxLength: 64<br>• minLength: 1 |
| interface.transport_address | string | query | False | Filter by interface.transport_address |
| interface.uuid | string | query | False | Filter by interface.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.name | string | query | False | Filter by interface.name |
| dh_hmac_chap.group_size | string | query | False | Filter by dh_hmac_chap.group_size<br>• Introduced in: 9.12 |
| dh_hmac_chap.hash_function | string | query | False | Filter by dh_hmac_chap.hash_function<br>• Introduced in: 9.12 |
| dh_hmac_chap.mode | string | query | False | Filter by dh_hmac_chap.mode<br>• Introduced in: 9.12 |
| digest.header | boolean | query | False | Filter by digest.header<br>• Introduced in: 9.15 |
| digest.data | boolean | query | False | Filter by digest.data<br>• Introduced in: 9.15 |
| id | string | query | False | Filter by id |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| io_queue.count | integer | query | False | Filter by io_queue.count |
| io_queue.depth | integer | query | False | Filter by io_queue.depth |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| keep_alive_timeout | integer | query | False | Filter by keep_alive_timeout<br>• Introduced in: 9.14 |
| tls.key_type | string | query | False | Filter by tls.key_type<br>• Introduced in: 9.16 |
| tls.cipher | string | query | False | Filter by tls.cipher<br>• Introduced in: 9.16 |
| tls.psk_identity | string | query | False | Filter by tls.psk_identity<br>• Introduced in: 9.16 |
| host.id | string | query | False | Filter by host.id |
| host.transport_address | string | query | False | Filter by host.transport_address |
| host.nqn | string | query | False | Filter by host.nqn<br>• maxLength: 223<br>• minLength: 1 |
| admin_queue.depth | integer | query | False | Filter by admin_queue.depth |
| transport_protocol | string | query | False | Filter by transport_protocol<br>• Introduced in: 9.16 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_controller] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "admin_queue": {
        "depth": 0
      },
      "dh_hmac_chap": {
        "group_size": "string",
        "hash_function": "string",
        "mode": "bidirectional"
      },
      "host": {
        "id": "b8546ca6097349e5b1558dc154fc073b",
        "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-
a3b0-71ce907b4eef",
        "transport_address": "nn-0x20000090fae00806:pn-
0x10000090fae00806"
      },
      "id": "0040h",
      "interface": {
        "name": "lif1",
        "transport_address": "nn-0x200400a0989a1c8d:pn-
0x200500a0989a1c8d",
        "uuid": "fa1c5941-2593-11e9-94c4-00a0989a1c8e"
      },
      "io_queue": {
        "count": 0,
        "depth": [
          "integer"
        ]
```

```
      },
      "keep_alive_timeout": 1500,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "subsystem": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tls": {
        "cipher": "tls_aes_128_gcm_sha256",
        "key_type": "configured",
        "psk_identity": "NVMe1R01 nqn.2014-
08.org.nvmexpress:uuid:713b3816-f9bf-ba43-b95a-5e4bf8c726e9 nqn.1992-
08.com.netapp:sn.76f9d9bfb96511eea95e005056bb72b2:subsystem.ss1
mS1A7nrooevA9ZqAM09fQzWQlB2UZRt0BE1X4vINjY0=:"
      },
      "transport_protocol": "string"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

admin_queue

| Name | Type | Description |
|---|---|---|
| depth | integer | The depth of the admin queue for the controller. |

dh_hmac_chap

A container for properties of the NVMe in-band authentication DH-HMAC-CHAP protocol used by the the host connection to the controller.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size used for NVMe in-band authentication. This property is populated only when NVMe in-band authentication was performed for the NVMe-oF transport connection. |
| hash_function | string | The hash function used for NVMe in-band authentication. This property is populated only when NVMe in-band authentication was performed for the NVMe-oF transport connection. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The NVMe in-band authentication mode used for the host connection. When set to:<br>• none: Neither the host nor controller was authenticated.<br>• unidirectional: The controller authenticated the host.<br>• bidirectional: The controller authenticated the host and the host authenticated the controller. |

digest

Digests are properties of NVMe controllers created over the NVMe/TCP transport protocol. The usage of digests is negotiated between the host and the controller during connection setup. ONTAP enables digests only if the host requests them. The header digest is the crc32 checksum of the header portion of the NVMe/TCP PDU. The data digest is the crc32 checksum of the data portion of the NVMe/TCP PDU.

If a digest is enabled, upon receiving an NVMe/TCP PDU, ONTAP calculates the crc32 checksum of the associated portion of the PDU and compares it with the digest value present in the transmitted PDU. If there is a mismatch, ONTAP returns an error and destroys the controller.

| Name | Type | Description |
| --- | --- | --- |
| data | boolean | Reports if digests are enabled for the data portion of the PDU. |
| header | boolean | Reports if digests are enabled for the header portion of the PDU. |

host

Properties of the connected host.

| Name | Type | Description |
| --- | --- | --- |
| id | string | The host identifier registered with the controller. |
| nqn | string | The NVMe qualified name of the host. |
| transport_address | string | The transport address of the host. |

interface

The logical interface through which the host is connected.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the logical interface. |
| transport_address | string | The transport address of the logical interface. |
| uuid | string | The unique identifier of the logical interface. |

io_queue

Properties of the I/O queues available to the controller.

| Name | Type | Description |
|---|---|---|
| count | integer | The number of I/O queues available to the controller. |
| depth | array[integer] | The depths of the I/O queues. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tls

A container for properties that describe the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| cipher | string | The cipher suite used for the transport by the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. This property is populated only when encryption is in use for the transport connection. |
| key_type | string | The method by which the TLS pre-shared key (PSK) was obtained when establishing the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem.<br>Possible values:<br>• `none` - TLS encryption is not configured for the host connection.<br>• `configured` - A user supplied PSK was used for the encrypted NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| psk_identity | string | The TLS PSK identity supplied by the host when establishing the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. This property is populated only when encryption is in use for the transport connection. |

nvme_subsystem_controller

A Non-Volatile Memory Express (NVMe) subsystem controller represents a connection between a host and a storage solution.

An NVMe subsystem controller is identified by the NVMe subsystem UUID and the controller ID.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| admin_queue | admin_queue | |
| dh_hmac_chap | dh_hmac_chap | A container for properties of the NVMe in-band authentication DH-HMAC-CHAP protocol used by the the host connection to the controller. |

| Name | Type | Description |
| --- | --- | --- |
| digest | digest | Digests are properties of NVMe controllers created over the NVMe/TCP transport protocol. The usage of digests is negotiated between the host and the controller during connection setup. ONTAP enables digests only if the host requests them. The header digest is the crc32 checksum of the header portion of the NVMe/TCP PDU. The data digest is the crc32 checksum of the data portion of the NVMe/TCP PDU.<br><br>If a digest is enabled, upon receiving an NVMe/TCP PDU, ONTAP calculates the crc32 checksum of the associated portion of the PDU and compares it with the digest value present in the transmitted PDU. If there is a mismatch, ONTAP returns an error and destroys the controller.<br><br>• readOnly: 1<br>• Introduced in: 9.15 |
| host | host | Properties of the connected host. |
| id | string | The identifier of the subsystem controller. This field consists of 4 zero-filled hexadecimal digits followed by an 'h'. |
| interface | interface | The logical interface through which the host is connected. |
| io_queue | io_queue | Properties of the I/O queues available to the controller. |

| Name | Type | Description |
| --- | --- | --- |
| keep_alive_timeout | integer | The keep-alive timeout value for the controller and all of its host connections, in milliseconds.<br><br>If the NVMe controller does not receive a keep-alive request or an I/O request within the timeout window, the NVMe controller terminates its admin queue and I/O queue connections leading to NVMe controller teardown. If the NVMe host does not receive a response to a keep-alive request or an I/O request within the timeout window, the NVMe host initiates a connection disconnect. |
| node | node | |
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tls | tls | A container for properties that describe the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. |
| transport_protocol | string | Transport Protocol |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe subsystem controller

GET `/protocols/nvme/subsystem-controllers/{subsystem.uuid}/{id}`

**Introduced In:** 9.6

Retrieves an NVMe subsystem controller.

#### Related ONTAP commands

- `vserver nvme subsystem controller show`

#### Learn more

- DOC /protocols/nvme/subsystem-controllers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| id | string | path | True | The unique identifier of the NVMe subsystem controller. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| admin_queue | admin_queue | |
| dh_hmac_chap | dh_hmac_chap | A container for properties of the NVMe in-band authentication DH-HMAC-CHAP protocol used by the the host connection to the controller. |
| digest | digest | Digests are properties of NVMe controllers created over the NVMe/TCP transport protocol. The usage of digests is negotiated between the host and the controller during connection setup. ONTAP enables digests only if the host requests them. The header digest is the crc32 checksum of the header portion of the NVMe/TCP PDU. The data digest is the crc32 checksum of the data portion of the NVMe/TCP PDU.<br><br>If a digest is enabled, upon receiving an NVMe/TCP PDU, ONTAP calculates the crc32 checksum of the associated portion of the PDU and compares it with the digest value present in the transmitted PDU. If there is a mismatch, ONTAP returns an error and destroys the controller.<br><br>• readOnly: 1<br>• Introduced in: 9.15 |
| host | host | Properties of the connected host. |
| id | string | The identifier of the subsystem controller. This field consists of 4 zero-filled hexadecimal digits followed by an 'h'. |
| interface | interface | The logical interface through which the host is connected. |
| io_queue | io_queue | Properties of the I/O queues available to the controller. |

| Name | Type | Description |
| --- | --- | --- |
| keep_alive_timeout | integer | The keep-alive timeout value for the controller and all of its host connections, in milliseconds.<br><br>If the NVMe controller does not receive a keep-alive request or an I/O request within the timeout window, the NVMe controller terminates its admin queue and I/O queue connections leading to NVMe controller teardown. If the NVMe host does not receive a response to a keep-alive request or an I/O request within the timeout window, the NVMe host initiates a connection disconnect. |
| node | node | |
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| tls | tls | A container for properties that describe the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. |
| transport_protocol | string | Transport Protocol |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "admin_queue": {
    "depth": 0
  },
  "dh_hmac_chap": {
    "group_size": "string",
    "hash_function": "string",
    "mode": "bidirectional"
  },
  "host": {
    "id": "b8546ca6097349e5b1558dc154fc073b",
    "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-a3b0-
71ce907b4eef",
    "transport_address": "nn-0x20000090fae00806:pn-0x10000090fae00806"
  },
  "id": "0040h",
  "interface": {
    "name": "lif1",
    "transport_address": "nn-0x200400a0989a1c8d:pn-0x200500a0989a1c8d",
    "uuid": "fa1c5941-2593-11e9-94c4-00a0989a1c8e"
  },
  "io_queue": {
    "count": 0,
    "depth": [
      "integer"
    ]
  },
  "keep_alive_timeout": 1500,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "subsystem": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tls": {
    "cipher": "tls_aes_128_gcm_sha256",
    "key_type": "configured",
    "psk_identity": "NVMe1R01 nqn.2014-08.org.nvmexpress:uuid:713b3816-
f9bf-ba43-b95a-5e4bf8c726e9 nqn.1992-
08.com.netapp:sn.76f9d9bfb96511eea95e005056bb72b2:subsystem.ss1
mS1A7nrooevA9ZqAM09fQzWQlB2UZRt0BE1X4vINjY0=:"
  },
  "transport_protocol": "string"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The supplied subsystem identifier does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

admin_queue

| Name | Type | Description |
|---|---|---|
| depth | integer | The depth of the admin queue for the controller. |

dh_hmac_chap

A container for properties of the NVMe in-band authentication DH-HMAC-CHAP protocol used by the the host connection to the controller.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size used for NVMe in-band authentication. This property is populated only when NVMe in-band authentication was performed for the NVMe-oF transport connection. |
| hash_function | string | The hash function used for NVMe in-band authentication. This property is populated only when NVMe in-band authentication was performed for the NVMe-oF transport connection. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The NVMe in-band authentication mode used for the host connection. When set to:<br>• none: Neither the host nor controller was authenticated.<br>• unidirectional: The controller authenticated the host.<br>• bidirectional: The controller authenticated the host and the host authenticated the controller. |

digest

Digests are properties of NVMe controllers created over the NVMe/TCP transport protocol. The usage of digests is negotiated between the host and the controller during connection setup. ONTAP enables digests only if the host requests them. The header digest is the crc32 checksum of the header portion of the NVMe/TCP PDU. The data digest is the crc32 checksum of the data portion of the NVMe/TCP PDU.

If a digest is enabled, upon receiving an NVMe/TCP PDU, ONTAP calculates the crc32 checksum of the associated portion of the PDU and compares it with the digest value present in the transmitted PDU. If there is a mismatch, ONTAP returns an error and destroys the controller.

| Name | Type | Description |
| --- | --- | --- |
| data | boolean | Reports if digests are enabled for the data portion of the PDU. |
| header | boolean | Reports if digests are enabled for the header portion of the PDU. |

host

Properties of the connected host.

| Name | Type | Description |
| --- | --- | --- |
| id | string | The host identifier registered with the controller. |
| nqn | string | The NVMe qualified name of the host. |
| transport_address | string | The transport address of the host. |

interface

The logical interface through which the host is connected.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the logical interface. |
| transport_address | string | The transport address of the logical interface. |
| uuid | string | The unique identifier of the logical interface. |

io_queue

Properties of the I/O queues available to the controller.

| Name | Type | Description |
|---|---|---|
| count | integer | The number of I/O queues available to the controller. |
| depth | array[integer] | The depths of the I/O queues. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tls

A container for properties that describe the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| cipher | string | The cipher suite used for the transport by the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. This property is populated only when encryption is in use for the transport connection. |
| key_type | string | The method by which the TLS pre-shared key (PSK) was obtained when establishing the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem.<br>Possible values:<br>• `none` - TLS encryption is not configured for the host connection.<br>• `configured` - A user supplied PSK was used for the encrypted NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| psk_identity | string | The TLS PSK identity supplied by the host when establishing the encrypted NVMe/TCP transport connection between the host and the NVMe subsystem. This property is populated only when encryption is in use for the transport connection. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NVMe subsystem maps

### Manage NVMe subsystem maps

#### Overview

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

The NVMe subsystem map REST API allows you to create, delete and discover NVMe subsystem maps.

#### Examples

#### Creating an NVMe subsystem map

```
# The API:
POST /api/protocols/nvme/subsystem-maps

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps' -H
'Accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"subsystem": { "name": "subsystem1" }, "namespace": { "name":
"/vol/vol1/namespace1" } }'
```

#### Retrieving all of the NVMe subsystem maps

```
# The API:
GET /api/protocols/nvme/subsystem-maps

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "0e91b214-fe40-11e8-91a0-005056a79967",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0e91b214-fe40-11e8-91a0-005056a79967"
        }
      }
    },
    "subsystem": {
      "uuid": "580a6b1e-fe43-11e8-91a0-005056a79967",
      "name": "subsystem1",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/580a6b1e-fe43-11e8-91a0-
005056a79967"
        }
      }
    },
```

```
    "namespace": {
      "uuid": "3ccdedc6-2519-4206-bc1f-b0f4adab6f89",
      "name": "/vol/vol1/namespace1",
      "_links": {
        "self": {
          "href": "/api/storage/namespaces/3ccdedc6-2519-4206-bc1f-
b0f4adab6f89"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-maps/580a6b1e-fe43-11e8-
91a0-005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystem-maps"
  }
}
}
```

#### Retrieving a specific NVMe subsystem map

The NVMe subsystem map is identified by the UUID of the NVMe subsystem followed by the UUID of the NVMe namespace.

```
# The API:
GET /api/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps/580a6b1e-
fe43-11e8-91a0-005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89' -H
'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "0e91b214-fe40-11e8-91a0-005056a79967",
  "name": "svm1",
```

```
  "_links": {
    "self": {
      "href": "/api/svm/svms/0e91b214-fe40-11e8-91a0-005056a79967"
    }
  }
},
"subsystem": {
  "uuid": "580a6b1e-fe43-11e8-91a0-005056a79967",
  "name": "subsystem1",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/580a6b1e-fe43-11e8-91a0-
005056a79967"
    }
  }
},
"namespace": {
  "uuid": "3ccdedc6-2519-4206-bc1f-b0f4adab6f89",
  "name": "/vol/vol1/namespace1",
  "node": {
    "name": "node1",
    "uuid": "012b4508-67d6-4788-8c2d-801f254ce976",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/012b4508-67d6-4788-8c2d-801f254ce976"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/namespaces/3ccdedc6-2519-4206-bc1f-
b0f4adab6f89"
    }
  }
},
"nsid": "00000001h",
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystem-maps/580a6b1e-fe43-11e8-91a0-
005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
  }
}
}
```

### Deleting an NVMe subsystem map

```
# The API:
DELETE /api/protocols/nvme/subsystem-
maps/{subsystem.uuid}/{namespace.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystem-
maps/580a6b1e-fe43-11e8-91a0-005056a79967/3ccdedc6-2519-4206-bc1f-
b0f4adab6f89' -H 'Accept: application/hal+json'
```

### Retrieve NVMe subsystem maps

GET `/protocols/nvme/subsystem-maps`

**Introduced In:** 9.6

Retrieves NVMe subsystem maps.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `anagrpid`

#### Related ONTAP commands

- `vserver nvme subsystem map show`

#### Learn more

- DOC /protocols/nvme/subsystem-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| subsystem.uuid | string | query | False | Filter by subsystem.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| subsystem.name | string | query | False | Filter by subsystem.name<br>• maxLength: 64<br>• minLength: 1 |
| nsid | string | query | False | Filter by nsid |
| namespace.node.name | string | query | False | Filter by namespace.node.name |
| namespace.node.uuid | string | query | False | Filter by namespace.node.uuid |
| namespace.uuid | string | query | False | Filter by namespace.uuid |
| namespace.name | string | query | False | Filter by namespace.name |
| anagrpid | string | query | False | Filter by anagrpid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_map] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpid": "00103050h",
      "namespace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/vol1/namespace1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h",
      "subsystem": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "subsystem1",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the NVMe namespace. Valid in POST.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the NVMe namespace. Valid in POST. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem_map

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| namespace | namespace | The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an NVMe subsystem map

POST `/protocols/nvme/subsystem-maps`

**Introduced In:** 9.6

Creates an NVMe subsystem map.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe subsystem map.
- `namespace.uuid` or `namespace.name` - Existing NVMe namespace to map to the specified NVme subsystem.
- `subsystem.uuid` or `subsystem.name` - Existing NVMe subsystem to map to the specified NVMe namespace.

#### Related ONTAP commands

- `vserver nvme subsystem map add`

#### Learn more

- DOC /protocols/nvme/subsystem-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| namespace | namespace | The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "anagrpid": "00103050h",
  "namespace": {
    "name": "/vol/vol1/namespace1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "nsid": "00000001h",
  "subsystem": {
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_map] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "anagrpid": "00103050h",
      "namespace": {
        "name": "/vol/vol1/namespace1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h",
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374127 | The specified namespace name is invalid. |

| Error Code | Description |
| --- | --- |
| 5376461 | The specified namespace name is invalid. |
| 5376462 | The specified namespace name is too long. |
| 72089731 | The specified NVMe namespace does not exist. |
| 72089790 | The supplied NVMe namespace is already mapped to the supplied NVMe subsystem. |
| 72089793 | An NVMe namespace in a snapshot cannot be mapped. |
| 72089799 | The NVMe namespace is the destination of an ongoing restore operation and is inaccessible for I/O and management. |
| 72089902 | A node does not have an NVMe interface configured. |
| 72089903 | Multiple nodes do not have an NVMe interface configured. |
| 72089904 | The aggregate must be given back to its home node prior to mapping the NVMe namespace it contains. |
| 72090001 | The NVMe subsystem specified by `subsystem.uuid` was not found. |
| 72090005 | The specified `namespace.uuid` and `namespace.name` refer to different NVMe namespaces. |
| 72090006 | The NVMe namespace specified by `namespace.uuid` was not found. |
| 72090007 | The NVMe namespace specified by `namespace.name` was not found. |
| 72090020 | The specified `subsystem.uuid` and `subsystem.name` refer to different NVMe subsystems. |
| 72090021 | The NVMe subsystem specified by `subsystem.name` was not found. |
| 72091102 | The NVMe namespace is part of a SnapMirror active sync relationship that has a different SVM peering relationship than the subsystem. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the NVMe namespace. Valid in POST.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | |
| uuid | string | The unique identifier of the NVMe namespace. Valid in POST. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem_map

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| namespace | namespace | The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both. |

| Name | Type | Description |
| --- | --- | --- |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an NVMe subsystem map**

DELETE `/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}`

**Introduced In:** 9.6

Deletes an NVMe subsystem map.

#### Related ONTAP commands

- `vserver nvme subsystem map remove`

#### Learn more

- DOC /protocols/nvme/subsystem-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| namespace.uuid | string | path | True | The unique identifier of the NVMe namespace. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_map] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpid": "00103050h",
      "namespace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/vol1/namespace1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h",
      "subsystem": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "subsystem1",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090019 | The specified NVMe namespace is not mapped to the specified NVMe subsystem. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the NVMe namespace. Valid in POST.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the NVMe namespace. Valid in POST. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem_map

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| namespace | namespace | The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe subsystem map

GET `/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}`

**Introduced In:** 9.6

Retrieves an NVMe subsystem map.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `anagrpid`

#### Related ONTAP commands

- `vserver nvme subsystem map show`

#### Learn more

- DOC /protocols/nvme/subsystem-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| namespace.uuid | string | path | True | The unique identifier of the NVMe namespace. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| namespace | namespace | The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "namespace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/vol1/namespace1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090019 | The specified NVMe namespace is not mapped to the specified NVMe subsystem. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the NVMe namespace. Valid in POST.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | |
| uuid | string | The unique identifier of the NVMe namespace. Valid in POST. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage NVMe subsystems

## Manage NVMe subsystems

### Overview

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

The NVMe subsystem REST API allows you to create, update, delete, and discover NVMe subsystems as well as add and remove NVMe hosts that can access the subsystem and associated namespaces.

### Examples

#### Creating an NVMe subsystem

```
# The API:
POST /api/protocols/nvme/subsystems

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'Accept:
application/json' -d '{ "svm": { "name": "svm1" }, "name": "subsystem1",
"os_type": "linux" }'
```

---

#### Creating an NVMe subsystem with multiple NVMe subsystem hosts

```
# The API:
POST /api/protocols/nvme/subsystems

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'Accept:
application/json' -d '{ "svm": { "name": "svm1" }, "name": "subsystem2",
"os_type": "vmware", "hosts": [ { "nqn": "nqn.1992-01.example.com:host1"
}, { "nqn": "nqn.1992-01.example.com:host2" } ] }'
```

---

#### Retrieving all NVMe subsystems

```
# The API:
GET /api/protocols/nvme/subsystems

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'Accept:
application/json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
    },
    "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
    "name": "subsystem1",
  },
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
    },
    "uuid": "bcde901a-a379-4a91-9ea6-1b728ed6696f",
    "name": "subsystem2",
  }
],
"num_records": 2,
}
```

---

#### Retrieving all NVMe subsystems with OS type *linux*

Note that the `os_type` query parameter is used to perform the query.

```
# The API:
GET /api/protocols/nvme/subsystems

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/nvme/subsystems?os_type=linux' -H 'Accept:
application/json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
    },
    "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
    "name": "subsystem1",
    "os_type": "linux",
  }
],
"num_records": 1,
}
```

### Retrieving a specific NVMe subsystem

```
# The API:
GET /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'Accept: application/json'

# The response:
{
"svm": {
  "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
  "name": "svm1",
},
"uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
"name": "subsystem1",
"os_type": "linux",
"target_nqn": "nqn.1992-
08.com.netapp:sn.d04594ef915b4c73b642169e72e4c0b1:subsystem.subsystem1",
"serial_number": "wtJNKNKD-uPLAAAAAAAD",
}
```

---

#### Retrieving the NVMe namespaces mapped to a specific NVMe subsystem

Note that the `fields` query parameter is used to specify the desired properties.

```
# The API:
GET /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?fields=subsystem_maps' -H 'Accept:
application/json'

# The response:
{
"svm": {
  "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
  "name": "svm1",
},
"uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
"name": "subsystem1",
"subsystem_maps": [
  {
    "anagrpid": "00000001h",
    "namespace": {
      "uuid": "eeaaca23-128d-4a7d-be4a-dc9106705799",
      "name": "/vol/vol1/namespace1"
    },
    "nsid": "00000001h"
  },
  {
    "anagrpid": "00000002h",
    "namespace": {
      "uuid": "feaaca23-83a0-4a7d-beda-dc9106705799",
      "name": "/vol/vol1/namespace2"
    },
    "nsid": "00000002h"
  }
]
}
```

### Adding a comment about an NVMe subsystem

```
# The API:
PATCH /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'Accept: application/json' -d '{
"comment": "A brief comment about the subsystem" }'
```

---

#### Deleting an NVMe subsystem

```
# The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'Accept: application/json'
```

#### Deleting an NVMe subsystem with mapped NVMe namespaces

Normally, deleting an NVMe subsystem that has mapped NVMe namespaces is not allowed. The deletion can be forced using the `allow_delete_while_mapped` query parameter.

```
# The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?allow_delete_while_mapped=true' -H 'Accept:
application/json'
```

#### Delete an NVMe subsystem with NVMe subsystem hosts

Normally, deleting an NVMe subsystem with NVMe subsystem hosts is disallowed. The deletion can be forced using the `allow_delete_with_hosts` query parameter.

```
# The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?allow_delete_with_hosts=true' -H 'Accept:
application/json'
```

---

#### An NVMe Subsystem Host

An NVMe subsystem host is a network host provisioned to an NVMe subsystem to access namespaces mapped to that subsystem.

#### Examples

### Adding an NVMe subsystem host to an NVMe subsystem

```
# The API:
POST /protocols/nvme/subsystems/{subsystem.uuid}/hosts

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'Accept: application/json' -d '{
"nqn": "nqn.1992-01.com.example:subsys1.host1" }'
```

---

### Adding multiple NVMe subsystem hosts to an NVMe subsystem

```
# The API:
POST /protocols/nvme/subsystems/{subsystem.uuid}/hosts

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'Accept: application/json' -d '{
"records": [ { "nqn": "nqn.1992-01.com.example:subsys1.host2" }, { "nqn":
"nqn.1992-01.com.example:subsys1.host3" } ] }'
```

---

### Retrieving all NVMe subsystem hosts for an NVMe subsystem

```
# The API:
GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'Accept: application/json'

# The response:
{
"records": [
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host1",
  },
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host2",
  },
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host3",
  }
],
"num_records": 3,
}
```

---

#### Retrieving a specific NVMe subsystem host for an NVMe subsystem

```
# The API:
GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example:subsys1.host1'
-H 'Accept: application/json'

# The response:
{
"subsystem": {
  "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
},
"nqn": "nqn.1992-01.com.example:subsys1.host1",
"dh_hmac_chap": {
  "node": "none"
},
"priority": "regular",
}
```

---

#### Deleting an NVMe subsystem host from an NVMe subsystem

```
# The API:
DELETE /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example:subsys1.host1'
-H 'Accept: application/json'
```

#### Retrieve NVMe subsystems

GET `/protocols/nvme/subsystems`

**Introduced In:** 9.6

Retrieves NVMe subsystems.

##### Related ONTAP commands

- `vserver nvme subsystem host show`
- `vserver nvme subsystem map show`
- `vserver nvme subsystem show`

**Learn more**

- DOC /protocols/nvme/subsystems

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| replication.peer_subsystem.uuid | string | query | False | Filter by replication.peer_subsystem.uuid<br>• Introduced in: 9.17 |
| replication.error.subsystem.local_svm | boolean | query | False | Filter by replication.error.subsystem.local_svm<br>• Introduced in: 9.17 |
| replication.error.subsystem.uuid | string | query | False | Filter by replication.error.subsystem.uuid<br>• Introduced in: 9.17 |
| replication.error.subsystem.name | string | query | False | Filter by replication.error.subsystem.name<br>• Introduced in: 9.17<br>• maxLength: 64<br>• minLength: 1 |
| replication.error.summary.arguments.message | string | query | False | Filter by replication.error.summary.arguments.message<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| replication.error.summary.arguments.code | string | query | False | Filter by replication.error.summary.arguments.code<br>• Introduced in: 9.17 |
| replication.error.summary.code | string | query | False | Filter by replication.error.summary.code<br>• Introduced in: 9.17 |
| replication.error.summary.message | string | query | False | Filter by replication.error.summary.message<br>• Introduced in: 9.17 |
| replication.state | string | query | False | Filter by replication.state<br>• Introduced in: 9.17 |
| replication.peer_svm.name | string | query | False | Filter by replication.peer_svm.name<br>• Introduced in: 9.17 |
| replication.peer_svm.uuid | string | query | False | Filter by replication.peer_svm.uuid<br>• Introduced in: 9.17 |
| serial_number | string | query | False | Filter by serial_number<br>• maxLength: 20<br>• minLength: 20 |
| svm.name | string | query | False | Filter by svm.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name<br>• maxLength: 64<br>• minLength: 1 |
| hosts.proximity.peer_svms.name | string | query | False | Filter by hosts.proximity.peer_svms.name<br>• Introduced in: 9.17 |
| hosts.proximity.peer_svms.uuid | string | query | False | Filter by hosts.proximity.peer_svms.uuid<br>• Introduced in: 9.17 |
| hosts.proximity.local_svm | boolean | query | False | Filter by hosts.proximity.local_svm<br>• Introduced in: 9.17 |
| hosts.dh_hmac_chap.group_size | string | query | False | Filter by hosts.dh_hmac_chap.group_size<br>• Introduced in: 9.12 |
| hosts.dh_hmac_chap.hash_function | string | query | False | Filter by hosts.dh_hmac_chap.hash_function<br>• Introduced in: 9.12 |
| hosts.dh_hmac_chap.mode | string | query | False | Filter by hosts.dh_hmac_chap.mode<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| hosts.nqn | string | query | False | Filter by hosts.nqn |
| hosts.priority | string | query | False | Filter by hosts.priority<br>• Introduced in: 9.14 |
| hosts.tls.key_type | string | query | False | Filter by hosts.tls.key_type<br>• Introduced in: 9.16 |
| io_queue.default.depth | integer | query | False | Filter by io_queue.default.depth<br>• Max value: 128<br>• Min value: 16 |
| io_queue.default.count | integer | query | False | Filter by io_queue.default.count<br>• Max value: 15<br>• Min value: 1 |
| delete_on_unmap | boolean | query | False | Filter by delete_on_unmap<br>• Introduced in: 9.7 |
| os_type | string | query | False | Filter by os_type |
| vendor_uuids | string | query | False | Filter by vendor_uuids<br>• Introduced in: 9.9 |
| comment | string | query | False | Filter by comment<br>• maxLength: 255<br>• minLength: 0 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | query | False | Filter by uuid |
| target_nqn | string | query | False | Filter by target_nqn<br>• maxLength: 223<br>• minLength: 1 |
| subsystem_maps.anagrpid | string | query | False | Filter by subsystem_maps.anagrpid |
| subsystem_maps.namespace.uuid | string | query | False | Filter by subsystem_maps.namespace.uuid |
| subsystem_maps.namespace.name | string | query | False | Filter by subsystem_maps.namespace.name |
| subsystem_maps.nsid | string | query | False | Filter by subsystem_maps.nsid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "hosts": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "dh_hmac_chap": {
            "group_size": "string",
            "hash_function": "string",
            "mode": "bidirectional"
          },
          "nqn": "nqn.1992-01.example.com:string",
          "priority": "string",
          "proximity": {
            "peer_svms": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "peer1",
                "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
              }
            ]
```

```
          },
          "tls": {
            "key_type": "configured"
          }
        }
      ],
      "io_queue": {
        "default": {
          "count": 4,
          "depth": 16
        }
      },
      "name": "subsystem1",
      "os_type": "string",
      "replication": {
        "error": {
          "subsystem": {
            "name": "subsystem1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "summary": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          }
        },
        "peer_subsystem": {
          "uuid": "1cd8a443-86d2-11e0-ae1c-123478563412"
        },
        "peer_svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "state": "string"
      },
      "serial_number": "wCVsgFMiuMhVAAAAAAAB",
```

```
      "subsystem_maps": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "anagrpid": "00103050h",
          "namespace": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "/vol/vol1/namespace1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "nsid": "00000001h"
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target_nqn": "nqn.1992-01.example.com:string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vendor_uuids": [
        "1447f0f4-42e5-0dfc-871a-dc9b3f92d8f8"
      ]
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies user configured values of I/O queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

default

Support for this field will be removed in a future release.
The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of host I/O queue pairs. |
| depth | integer | The host I/O queue depth. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
|---|---|---|
| default | default | Support for this field will be removed in a future release. The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem. |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | Indicates whether the reported subsystem is on the local SVM or the peer SVM. When deleting a replicated subsystem, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer subsystem remains, the peer subsystem is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the error.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
|---|---|---|
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |
| summary | summary | A user friendly message describing the error. |

peer_subsystem

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the peer subsystem. |

peer_svm

The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to subsystem replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_subsystem | peer_subsystem | |

| Name | Type | Description |
| --- | --- | --- |
| peer_svm | peer_svm | The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group. |
| state | string | The state of the replication queue associated with this subsystem. If this subsystem is not in the replication queue, the state is reported as *ok*. If this subsystem is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this subsystem is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

namespace

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

subsystem_maps

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| namespace | namespace | An NVMe namespace mapped to the NVMe subsystem. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |
| vendor_uuids | array[string] | Vendor-specific identifiers (UUIDs) optionally assigned to an NVMe subsystem when the subsystem is created. The identifiers are used to enable vendor-specific NVMe protocol features. The identifiers are provided by a host application vendor and shared with NetApp prior to a joint product release. Creating an NVMe subsystem with an unknown or non-specific identifier will have no effect on the NVMe subsystem. Refer to the ONTAP SAN Administration Guide for a list of the supported vendor-specific identifiers. After a subsystem is created, the vendor-specific identifiers cannot be changed or removed. Optional in POST.<br>• Introduced in: 9.9<br>• readCreate: 1 |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an NVMe subsystem

POST `/protocols/nvme/subsystems`

**Introduced In:** 9.6

Creates an NVMe subsystem.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe subsystem.
- `name` - Name for NVMe subsystem. Once created, an NVMe subsystem cannot be renamed.
- `os_type` - Operating system of the NVMe subsystem's hosts.

#### Related ONTAP commands

- `vserver nvme subsystem create`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |

| Name | Type | Description |
| --- | --- | --- |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| vendor_uuids | array[string] | Vendor-specific identifiers (UUIDs) optionally assigned to an NVMe subsystem when the subsystem is created. The identifiers are used to enable vendor-specific NVMe protocol features. The identifiers are provided by a host application vendor and shared with NetApp prior to a joint product release. Creating an NVMe subsystem with an unknown or non-specific identifier will have no effect on the NVMe subsystem. Refer to the ONTAP SAN Administration Guide for a list of the supported vendor-specific identifiers. After a subsystem is created, the vendor-specific identifiers cannot be changed or removed. Optional in POST.<br>• Introduced in: 9.9<br>• readCreate: 1 |

**Example request**


```
{
  "comment": "string",
  "hosts": [
    {
      "dh_hmac_chap": {
        "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "group_size": "string",
        "hash_function": "string",
        "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "mode": "bidirectional"
      },
      "nqn": "nqn.1992-01.example.com:string",
      "priority": "string",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      },
      "tls": {
        "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
        "key_type": "configured"
      }
    }
  ],
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
  },
  "name": "subsystem1",
  "os_type": "string",
  "replication": {
    "error": {
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
```

```
      },
      "peer_subsystem": {
        "uuid": "1cd8a443-86d2-11e0-ae1c-123478563412"
      },
      "peer_svm": {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      },
      "state": "string"
    },
    "serial_number": "wCVsgFMiuMhVAAAAAAAB",
    "subsystem_maps": [
      {
        "anagrpid": "00103050h",
        "namespace": {
          "name": "/vol/vol1/namespace1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "nsid": "00000001h"
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "target_nqn": "nqn.1992-01.example.com:string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "vendor_uuids": [
      "1447f0f4-42e5-0dfc-871a-dc9b3f92d8f8"
    ]
  }
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "comment": "string",
      "hosts": [
        {
          "dh_hmac_chap": {
            "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "group_size": "string",
            "hash_function": "string",
            "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "mode": "bidirectional"
          },
          "nqn": "nqn.1992-01.example.com:string",
          "priority": "string",
          "proximity": {
            "peer_svms": [
              {
                "name": "peer1",
                "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
              }
            ]
          },
          "tls": {
            "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
            "key_type": "configured"
          }
        }
      ],
      "io_queue": {
        "default": {
          "count": 4,
          "depth": 16
        }
      },
      "name": "subsystem1",
      "os_type": "string",
      "replication": {
        "error": {
          "subsystem": {
```

```
            "name": "subsystem1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "peer_subsystem": {
          "uuid": "1cd8a443-86d2-11e0-ae1c-123478563412"
        },
        "peer_svm": {
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "state": "string"
      },
      "serial_number": "wCVsgFMiuMhVAAAAAAAB",
      "subsystem_maps": [
        {
          "anagrpid": "00103050h",
          "namespace": {
            "name": "/vol/vol1/namespace1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "nsid": "00000001h"
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target_nqn": "nqn.1992-01.example.com:string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vendor_uuids": [
        "1447f0f4-42e5-0dfc-871a-dc9b3f92d8f8"
      ]
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The supplied SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | The svm.uuid or svm.name must be provided. |
| 72089635 | Setting vendor-specific UUIDs on NVMe subsystems is not supported until the effective cluster version is 9.9 or later. |
| 72089636 | Creating NVMe subsystems with `os_type` AIX is not supported until the effective cluster version is 9.13.1 or later. |
| 72089709 | The NVMe subsystem name contains an invalid character. |
| 72089711 | An invalid vendor-specific UUID was specified. |
| 72089712 | A duplicate vendor-specific UUID was specific. |
| 72089713 | Too many vendor UUIDs were supplied. |
| 72089716 | The DH-HMAC-CHAP secret property is invalid. DH-HMAC-CHAP secrets must be in the format "DHHC-1:0X:<Base 64 encoded key and CRC>:", where X represents 0, 1, or 3 indicating no hash function, SHA-256, and SHA-512 respectively. |
| 72089771 | The NQN is invalid. A non-empty qualifier is required after the prefix. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089772 | The NQN is invalid. Add the prefix *'nqn'*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089773 | The NQN is invalid. The date field must be formatted *yyyy-mm*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72090003 | A host to be added to an NVMe subsystem is missing the "nqn" property. |
| 72090025 | The NVMe subsystem already exists for the SVM. |
| 72090029 | The NVMe service does not exist. |
| 72090030 | A partial success occurred while adding multiple NVMe subsystem hosts to an NVMe subsystem. |

| Error Code | Description |
| --- | --- |
| 72090036 | An NVMe subsystem host NQN was duplicated in the input. |
| 72090042 | The DH-HMAC-CHAP secret property is required when setting any other NVMe in-band authentication properties for a host. |
| 72090043 | An igroup already exists with the requested NVMe subsystem name. |
| 72090151 | NVMe/TCP-TLS is not supported for the effective version of the cluster. |
| 72090202 | A provided NVMe subsystem host TLS configured PSK is not valid. |
| 72090204 | A TLS configured PSK was not provided when adding an NVMe subsystem host with the configured key type. |
| 72090205 | An invalid combination for the TLS key type and configured PSK values was provided when adding an NVMe subsystem host. When key type is "none", no configured PSK is allowed. When key type is "configured", a configured PSK is required. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| controller_secret_key | string | The controller secret for NVMe in-band authentication. The value of this property is used by the NVMe host to authenticate the NVMe controller while establishing a connection. If unset, the controller is not authenticated. When supplied, the property `host_secret_key` must also be supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a controller secret has been set for the host, but the controller secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| host_secret_key | string | The host secret for NVMe in-band authentication. The value of this property is used by the NVMe controller to authenticate the NVMe host while establishing a connection. If unset, no authentication is performed by the host or controller. This property must be supplied if any other NVMe in-band authentication properties are supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a host secret has been set for the host, but the host secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |

| Name | Type | Description |
| --- | --- | --- |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| configured_psk | string | A user supplied pre-shared key (PSK) value in PSK Interchange Format. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree. This property is only allowed when `key_type` is `configured`. If `configured_psk` is supplied and `key_type` is unset, `key_type` defaults to `configured`.<br><br>This property is write-only. The `key_type` property can be used to identify if a configured PSK has been set for the host, but the PSK value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies user configured values of I/O queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

default

Support for this field will be removed in a future release.
The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of host I/O queue pairs. |
| depth | integer | The host I/O queue depth. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| default | default | Support for this field will be removed in a future release. The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem. |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | Indicates whether the reported subsystem is on the local SVM or the peer SVM. When deleting a replicated subsystem, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer subsystem remains, the peer subsystem is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the error.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
| --- | --- | --- |
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |

peer_subsystem

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the peer subsystem. |

peer_svm

The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to subsystem replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_subsystem | peer_subsystem | |
| peer_svm | peer_svm | The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The state of the replication queue associated with this subsystem. If this subsystem is not in the replication queue, the state is reported as *ok*. If this subsystem is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this subsystem is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

namespace

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

subsystem_maps

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| namespace | namespace | An NVMe namespace mapped to the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| vendor_uuids | array[string] | Vendor-specific identifiers (UUIDs) optionally assigned to an NVMe subsystem when the subsystem is created. The identifiers are used to enable vendor-specific NVMe protocol features. The identifiers are provided by a host application vendor and shared with NetApp prior to a joint product release. Creating an NVMe subsystem with an unknown or non-specific identifier will have no effect on the NVMe subsystem. Refer to the ONTAP SAN Administration Guide for a list of the supported vendor-specific identifiers. After a subsystem is created, the vendor-specific identifiers cannot be changed or removed. Optional in POST.<br>• Introduced in: 9.9<br>• readCreate: 1 |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve NVMe subsystem hosts**

GET `/protocols/nvme/subsystems/{subsystem.uuid}/hosts`

**Introduced In:** 9.6

Retrieves the NVMe subsystem hosts of an NVMe subsystem.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `subsystem_maps.*`

#### Related ONTAP commands

- `vserver nvme subsystem map show`
- `vserver nvme subsystem show`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_host] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "dh_hmac_chap": {
        "group_size": "string",
        "hash_function": "string",
        "mode": "bidirectional"
      },
      "io_queue": {
        "count": 4,
        "depth": 32
      },
      "nqn": "nqn.1992-01.example.com:string",
      "priority": "string",
      "proximity": {
        "peer_svms": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      },
      "subsystem": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "tls": {
        "key_type": "configured"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The NVMe subsystem does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue count is being used. Valid in GET only. |
| depth | integer | The I/O queue depth. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue depth is being used. Valid in GET only. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem. |
| depth | integer | The I/O queue depth. The default value is inherited from the owning NVMe subsystem. |

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the `records` property is used. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

nvme_subsystem_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add NVMe subsystem hosts

POST `/protocols/nvme/subsystems/{subsystem.uuid}/hosts`

**Introduced In:** 9.6

Adds NVMe subsystem host(s) to an NVMe subsystem.

#### Required properties

- `nqn` or `records.nqn` - NVMe host(s) NQN(s) to add to the NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem host add`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |

| Name | Type | Description |
| --- | --- | --- |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| records | array[records] | An array of NVMe hosts specified to add multiple NVMe hosts to an NVMe subsystem in a single API call. Valid in POST only. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

**Example request**


```
{
  "dh_hmac_chap": {
    "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
    "group_size": "string",
    "hash_function": "string",
    "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
    "mode": "bidirectional"
  },
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "nqn": "nqn.1992-01.example.com:string",
  "priority": "string",
  "proximity": {
    "peer_svms": [
      {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
  },
  "records": [
    {
      "dh_hmac_chap": {
        "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "group_size": "string",
        "hash_function": "string",
        "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "mode": "bidirectional"
      },
      "io_queue": {
        "count": 4,
        "depth": 32
      },
      "nqn": "nqn.1992-01.example.com:string",
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
```

```
      "tls": {
        "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
        "key_type": "configured"
      }
    }
  ],
  "subsystem": {
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tls": {
    "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
    "key_type": "configured"
  }
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[nvme_subsystem_host] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "dh_hmac_chap": {
        "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "group_size": "string",
        "hash_function": "string",
        "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
        "mode": "bidirectional"
      },
      "io_queue": {
        "count": 4,
        "depth": 32
      },
      "nqn": "nqn.1992-01.example.com:string",
      "priority": "string",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      },
      "records": [
        {
          "dh_hmac_chap": {
            "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "group_size": "string",
            "hash_function": "string",
            "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "mode": "bidirectional"
          },
          "io_queue": {
            "count": 4,
            "depth": 32
          },
          "nqn": "nqn.1992-01.example.com:string",
          "subsystem": {
```

```
            "name": "subsystem1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "tls": {
            "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
            "key_type": "configured"
          }
        }
      ],
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "tls": {
        "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
        "key_type": "configured"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262186 | The "records" array and other host properties are mutually exclusive. |
| 72089705 | The NVMe subsystem host already exists for the NVMe subsystem. |

| Error Code | Description |
| --- | --- |
| 72089716 | The DH-HMAC-CHAP secret property is invalid. DH-HMAC-CHAP secrets must be in the format "DHHC-1:0X:<Base 64 encoded key and CRC>:", where X represents 0, 1, or 3 indicating no hash function, SHA-256, and SHA-512 respectively. |
| 72089771 | The NQN is invalid. A non-empty qualifier is required after the prefix. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089772 | The NQN is invalid. Add the prefix *'nqn'*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089773 | The NQN is invalid. The date field must be formatted *yyyy-mm*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72090001 | The NVMe subsystem does not exist. |
| 72090003 | A host to be added to an NVMe subsystem is missing the "nqn" property. |
| 72090036 | An NVMe subsystem host NQN is duplicated in the input. |
| 72090041 | An element in the "records" array contains an invalid property. |
| 72090042 | The DH-HMAC-CHAP secret property is required when setting any other NVMe in-band authentication properties for a host. |
| 72090151 | NVMe/TCP-TLS is not supported for the effective version of the cluster. |
| 72090202 | A provided NVMe subsystem host TLS configured PSK is not valid. |
| 72090204 | A TLS configured PSK was not provided when adding an NVMe subsystem host with the configured key type. |
| 72090205 | An invalid combination for the TLS key type and configured PSK values was provided when adding an NVMe subsystem host. When key type is "none", no configured PSK is allowed. When key type is "configured", a configured PSK is required. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| controller_secret_key | string | The controller secret for NVMe in-band authentication. The value of this property is used by the NVMe host to authenticate the NVMe controller while establishing a connection. If unset, the controller is not authenticated. When supplied, the property `host_secret_key` must also be supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a controller secret has been set for the host, but the controller secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| host_secret_key | string | The host secret for NVMe in-band authentication. The value of this property is used by the NVMe controller to authenticate the NVMe host while establishing a connection. If unset, no authentication is performed by the host or controller. This property must be supplied if any other NVMe in-band authentication properties are supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a host secret has been set for the host, but the host secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue count is being used. Valid in GET only. |
| depth | integer | The I/O queue depth. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue depth is being used. Valid in GET only. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem. |
| depth | integer | The I/O queue depth. The default value is inherited from the owning NVMe subsystem. |

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| configured_psk | string | A user supplied pre-shared key (PSK) value in PSK Interchange Format. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree. This property is only allowed when `key_type` is `configured`. If `configured_psk` is supplied and `key_type` is unset, `key_type` defaults to `configured`.<br><br>This property is write-only. The `key_type` property can be used to identify if a configured PSK has been set for the host, but the PSK value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |

| Name | Type | Description |
| --- | --- | --- |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the `records` property is used. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

nvme_subsystem_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| records | array[records] | An array of NVMe hosts specified to add multiple NVMe hosts to an NVMe subsystem in a single API call. Valid in POST only. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete an NVMe subsystem host

DELETE `/protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}`

**Introduced In:** 9.6

Deletes an NVMe subsystem host from an NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem host remove`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| nqn | string | path | True | The NVMe qualified name (NQN) used to identify the NVMe subsystem host. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72089771 | The NQN is invalid. A non-empty qualifier is required after the prefix. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089772 | The NQN is invalid. Add the prefix *'nqn'*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72089773 | The NQN is invalid. The date field must be formatted *yyyy-mm*. An example of a valid NQN is *nqn.1992-01.com.example:string*. |
| 72090001 | The NVMe subsystem does not exist. |
| 72090004 | The NVMe subsystem host does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe subsystem host

GET `/protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}`

**Introduced In:** 9.6

Retrieves an NVMe subsystem host of an NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem host show`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nqn | string | path | True | The NVMe qualified name (NQN) used to identify the NVMe subsystem host. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "dh_hmac_chap": {
    "group_size": "string",
    "hash_function": "string",
    "mode": "bidirectional"
  },
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "nqn": "nqn.1992-01.example.com:string",
  "priority": "string",
  "proximity": {
    "peer_svms": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
  },
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tls": {
    "key_type": "configured"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The NVMe subsystem does not exist. |
| 72090004 | The NVMe subsystem host does not exist. |
| 72090022 | The NVMe subsystem host does not exist in the specified subsystem. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue count is being used. Valid in GET only. |
| depth | integer | The I/O queue depth. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue depth is being used. Valid in GET only. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem. |
| depth | integer | The I/O queue depth. The default value is inherited from the owning NVMe subsystem. |

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the `records` property is used. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an NVMe subsystem host

PATCH `/protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}`

**Introduced In:** 9.17

Updates an NVMe subsystem host.

#### Related ONTAP commands

- `vserver nvme subsystem host add-proximal-vserver`
- `vserver nvme subsystem host remove-proximal-vserver`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem.uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| nqn | string | path | True | The NVMe qualified name (NQN) used to identify the NVMe subsystem host. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |

**Example request**

```
{
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "proximity": {
    "peer_svms": [
      {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
  },
  "subsystem": {
    "name": "subsystem1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5376354 | The SVM peer relationship does not exist. |
| 26345675 | The SVM peer relationship contains local SVM name mismatch. |
| 72090001 | The NVMe subsystem does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
|---|---|---|
| count | integer | The number of I/O queue pairs. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue count is being used. Valid in GET only. |

| Name | Type | Description |
| --- | --- | --- |
| depth | integer | The I/O queue depth. Absence of this property in GET implies property priority is set and platform and transport protocol specific values for I/O queue depth is being used. Valid in GET only. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem. |
| depth | integer | The I/O queue depth. The default value is inherited from the owning NVMe subsystem. |

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |

nvme_subsystem_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| subsystem | subsystem | The NVMe subsystem to which the NVMe host has been provisioned. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Remove an NVMe subsystem

DELETE `/protocols/nvme/subsystems/{uuid}`

**Introduced In:** 9.6

Removes an NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem delete`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| allow_delete_while_mapped | boolean | query | False | Allows for the deletion of a mapped NVMe subsystem. |
| allow_delete_with_hosts | boolean | query | False | Allows for the deletion of an NVMe subsystem with NVMe hosts. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The NVMe subsystem does not exist. |

| Error Code | Description |
| --- | --- |
| 72090023 | The NVMe subsystem contains one or more mapped namespaces. Use the `allow_delete_while_mapped` query parameter to delete an NVMe subsystem with mapped NVMe namespaces. |
| 72090024 | The NVMe subsystem contains one or more NVMe hosts. Use the `allow_delete_with_hosts` query parameter to delete an NVMe subsystem with NVMe hosts. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe subsystem

GET `/protocols/nvme/subsystems/{uuid}`

**Introduced In:** 9.6

Retrieves an NVMe subsystem.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `subsystem_maps.*`

#### Related ONTAP commands

- `vserver nvme subsystem host show`
- `vserver nvme subsystem map show`
- `vserver nvme subsystem show`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe subsystem. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| vendor_uuids | array[string] | Vendor-specific identifiers (UUIDs) optionally assigned to an NVMe subsystem when the subsystem is created. The identifiers are used to enable vendor-specific NVMe protocol features. The identifiers are provided by a host application vendor and shared with NetApp prior to a joint product release. Creating an NVMe subsystem with an unknown or non-specific identifier will have no effect on the NVMe subsystem. Refer to the ONTAP SAN Administration Guide for a list of the supported vendor-specific identifiers. After a subsystem is created, the vendor-specific identifiers cannot be changed or removed. Optional in POST.<br>• Introduced in: 9.9<br>• readCreate: 1 |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "hosts": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "dh_hmac_chap": {
        "group_size": "string",
        "hash_function": "string",
        "mode": "bidirectional"
      },
      "nqn": "nqn.1992-01.example.com:string",
      "priority": "string",
      "proximity": {
        "peer_svms": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      },
      "tls": {
        "key_type": "configured"
      }
    }
  ],
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
```

```
  },
  "name": "subsystem1",
  "os_type": "string",
  "replication": {
    "error": {
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "summary": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    },
    "peer_subsystem": {
      "uuid": "1cd8a443-86d2-11e0-ae1c-123478563412"
    },
    "peer_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "peer1",
      "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
    },
    "state": "string"
  },
  "serial_number": "wCVsgFMiuMhVAAAAAAAB",
  "subsystem_maps": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpid": "00103050h",
      "namespace": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/vol1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target_nqn": "nqn.1992-01.example.com:string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vendor_uuids": [
    "1447f0f4-42e5-0dfc-871a-dc9b3f92d8f8"
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The NVMe subsystem does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies user configured values of I/O queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

default

Support for this field will be removed in a future release.
The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of host I/O queue pairs. |
| depth | integer | The host I/O queue depth. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
|---|---|---|
| default | default | Support for this field will be removed in a future release. The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem. |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | Indicates whether the reported subsystem is on the local SVM or the peer SVM. When deleting a replicated subsystem, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer subsystem remains, the peer subsystem is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the error.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
|---|---|---|
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |
| summary | summary | A user friendly message describing the error. |

peer_subsystem

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the peer subsystem. |

peer_svm

The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to subsystem replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_subsystem | peer_subsystem | |

| Name | Type | Description |
| --- | --- | --- |
| peer_svm | peer_svm | The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group. |
| state | string | The state of the replication queue associated with this subsystem. If this subsystem is not in the replication queue, the state is reported as *ok*. If this subsystem is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this subsystem is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

namespace

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

subsystem_maps

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| namespace | namespace | An NVMe namespace mapped to the NVMe subsystem. |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Update an NVMe subsystem

PATCH `/protocols/nvme/subsystems/{uuid}`

**Introduced In:** 9.6

Updates an NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem modify`

#### Learn more

- DOC /protocols/nvme/subsystems

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | The unique identifier of the NVMe subsystem. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |

| Name | Type | Description |
| --- | --- | --- |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |

### Example request


```
{
  "comment": "string",
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
  },
  "replication": {
    "error": {
      "subsystem": {
        "name": "subsystem1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "peer_subsystem": {
      "uuid": "1cd8a443-86d2-11e0-ae1c-123478563412"
    },
    "peer_svm": {
      "name": "peer1",
      "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
    },
    "state": "string"
  },
  "serial_number": "wCVsgFMiuMhVAAAAAAAB",
  "subsystem_maps": [
    {
      "anagrpid": "00103050h",
      "namespace": {
        "name": "/vol/vol1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h"
    }
  ],
  "target_nqn": "nqn.1992-01.example.com:string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090001 | The NVMe subsystem does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
|---|---|---|
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

hosts

default

> Support for this field will be removed in a future release.
> The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of host I/O queue pairs. |
| depth | integer | The host I/O queue depth. |

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

| Name | Type | Description |
| --- | --- | --- |
| default | default | Support for this field will be removed in a future release. The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem. |

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | Indicates whether the reported subsystem is on the local SVM or the peer SVM. When deleting a replicated subsystem, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer subsystem remains, the peer subsystem is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the NVMe subsystem. |
| uuid | string | The unique identifier of the NVMe subsystem. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the error.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
| --- | --- | --- |
| subsystem | subsystem | An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts. |

peer_subsystem

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the peer subsystem. |

peer_svm

The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to subsystem replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this subsystem. Subsystems within a peering relationship are replicated in the same stream, so the error reported here might be related to this subsystem or a prior replicated subsystem that is now blocking the replication of this subsystem. Both the error information and the subsystem encountering the error are reported. If the error is configuration related, it can be corrected on the referenced subsystem. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_subsystem | peer_subsystem | |
| peer_svm | peer_svm | The peered SVM to which the subsystem is replicated. Subsystem are are automatically replicated when mapped to a namespace in a SnapMirror active sync relationship. When a subsystem is mapped to a namespace in an active sync relationship, the subsystem is restricted to only be mapped to namespaces that are members of the same consistency group. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The state of the replication queue associated with this subsystem. If this subsystem is not in the replication queue, the state is reported as *ok*. If this subsystem is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this subsystem is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

namespace

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

subsystem_maps

An NVMe namespace mapped to the NVMe subsystem.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| namespace | namespace | An NVMe namespace mapped to the NVMe subsystem. |

| Name | Type | Description |
| --- | --- | --- |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST and PATCH. |
| delete_on_unmap | boolean | An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the subsystem is created. |
| io_queue | io_queue | The properties of the submission queue used to submit I/O commands for execution by the NVMe controller. |

| Name | Type | Description |
| --- | --- | --- |
| replication | replication | Properties related to subsystem replication. |
| serial_number | string | The serial number of the NVMe subsystem. |
| subsystem_maps | array[subsystem_maps] | The NVMe namespaces mapped to the NVMe subsystem.<br><br>There is an added computational cost to retrieving property values for `subsystem_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| target_nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |
| uuid | string | The unique identifier of the NVMe subsystem. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NVMe namespaces

### Manage NVMe namespaces

#### Overview

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

The NVMe namespace REST API allows you to create, update, delete and discover NVMe namespaces.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

#### Performance monitoring

Performance of an NVMe namespace can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the space utilization and performance of an NVMe namespace in terms of IOPS, latency, and throughput. The `metric.*` properties denote an average, whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

### Creating an NVMe namespace

This example creates a 300 gigabyte NVMe namespace, with 4096-byte blocks, in SVM *svm1*, volume *vol1*, configured for use by *linux* hosts. The `return_records` query parameter is used to retrieve properties of the newly created NVMe namespace in the POST response.

```
# The API:
POST /api/storage/namespaces

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/namespaces?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "os_type": "linux",
"space": { "block_size": "4096", "size": "300G" }, "name" :
"/vol/vol1/namespace1" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "dccdc3e6-cf4e-498f-bec6-f7897f945669",
    "svm": {
```

```
      "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
        }
      }
    },
    "name": "/vol/vol1/namespace1",
    "location": {
      "namespace": "namespace1",
      "volume": {
        "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-
005056bbc17d"
          }
        }
      }
    },
    "enabled": true,
    "os_type": "linux",
    "space": {
      "block_size": 4096,
      "size": 322122547200,
      "used": 0,
      "guarantee": {
        "requested": false,
        "reserved": false
      }
    },
    "status": {
      "container_state": "online",
      "read_only": false,
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669"
      }
    }
  }
]
```

```
}
```

#### Updating an NVMe namespace comment

This example sets the `comment` property of an NVMe namespace.

```
# The API:
PATCH /api/storage/namespaces/{uuid}

# The call:
```

#### Updating the size of an NVMe namespace

This example increases the size of an NVMe namespace.

```
# The API:
PATCH /api/storage/namespaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-
498f-bec6-f7897f945669' -H 'Accept: application/hal+json' -d '{ "space": {
"size": "1073741824" } }'
```

#### Retrieving NVMe namespaces

This example retrieves summary information for all online NVMe namespaces in SVM *svm1*. The `svm.name` and `status.state` query parameters are to find the desired NVMe namespaces.

```
# The API:
GET /api/storage/namespaces

# The call:
curl -X GET 'https://<mgmt-
ip>/api/storage/namespaces?svm.name=svm1&status.state=online' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "uuid": "5c254d22-96a6-42ac-aad8-0cd9ebd126b6",
```

```
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol1/namespace2",
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/namespaces/5c254d22-96a6-42ac-aad8-
0cd9ebd126b6"
      }
    }
  },
  {
    "uuid": "dccdc3e6-cf4e-498f-bec6-f7897f945669",
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol1/namespace1",
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669"
      }
    }
  },
  {
    "uuid": "be732687-20cf-47d2-a0e2-2a989d15661d",
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol2/namespace3",
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/namespaces/be732687-20cf-47d2-a0e2-
2a989d15661d"
      }
    }
  }
```

```
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/namespaces?svm.name=svm1&status.state=online"
  }
}
}
```

#### Retrieving details for a specific NVMe namespace

In this example, the `fields` query parameter is used to request all fields, including advanced fields, that would not otherwise be returned by default for the NVMe namespace.

```
# The API:
GET /api/storage/namespaces/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-498f-
bec6-f7897f945669?fields=**' -H 'Accept: application/hal+json'

# The response:
{
"uuid": "dccdc3e6-cf4e-498f-bec6-f7897f945669",
"svm": {
  "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
    }
  }
},
"name": "/vol/vol1/namespace1",
"location": {
  "namespace": "namespace1",
  "volume": {
    "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-
005056bbc17d"
      }
```

```
    }
  }
},
"auto_delete": false,
"enabled": true,
"comment": "Data for the research department.",
"os_type": "linux",
"space": {
  "block_size": 4096,
  "size": 322122547200,
  "used": 0,
  "guarantee": {
    "requested": false,
    "reserved": false
  }
},
"status": {
  "container_state": "online",
  "mapped": true,
  "read_only": false,
  "state": "online"
},
"subsystem_map": {
  "nsid": "00000001h",
  "anagrpid": "00000001h",
  "subsystem": {
    "uuid": "01f17d05-2be9-11e9-bed2-005056bbc17d",
    "name": "subsystem1",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/01f17d05-2be9-11e9-bed2-
005056bbc17d"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem-maps/dccdc3e6-cf4e-498f-bec6-
f7897f945669/01f17d05-2be9-11e9-bed2-005056bbc17d"
    }
  }
},
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "duration": "PT15S",
  "status": "ok",
```

```
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669?fields=**"
  }
}
}
```

**Cloning NVMe namespaces**

A clone of an NVMe namespace is an independent "copy" of the namespace that shares unchanged data blocks with the original. As blocks of the source and clone are modified, unique blocks are written for each. NVMe namespace clones can be created quickly and consume very little space initially. They can be created for the purpose of back-up, or to replicate data for multiple consumers.

An NVMe namespace clone can also be set to auto-delete by setting the `auto_delete` property. If the namespace's volume is configured for automatic deletion, NVMe namespaces that have auto-delete enabled are deleted when a volume is nearly full to reclaim a target amount of free space in the volume.

### Creating a new NVMe namespace clone

You create an NVMe namespace clone as you create any NVMe namespace — a POST to /storage/namespaces. Set `clone.source.uuid` or `clone.source.name` to identify the source NVMe namespace from which the clone is created. The NVMe namespace clone and its source must reside in the same volume.

The source NVMe namespace can reside in a snapshot, in which case, the `clone.source.name` field must be used to identify it. Add `/.snapshot/<snapshot_name>` to the path after the volume name to identify the snapshot. For example `/vol/vol1/.snapshot/snap1/namespace1`.

```
# The API:
POST /api/storage/namespaces

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/namespaces' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name":
"/vol/vol1/namespace2clone1", "clone": { "source": { "name":
"/vol/vol1/namespace2" } } }'
```

### Over-writing an existing NVMe namespace's data as a clone of another

You can over-write an existing NVMe namespace as a clone of another. You do this as a PATCH on the NVMe namespace to overwrite — a PATCH to /storage/namespaces/{uuid}. Set the `clone.source.uuid` or `clone.source.name` property to identify the source NVMe namespace from which the clone data is taken. The NVMe namespace clone and its source must reside in the same volume.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete`, `subsystem_map`, `status.state`, and `uuid`.

```
# The API:
PATCH /api/storage/namespaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-
498f-bec6-f7897f945669' -H 'Accept: application/hal+json' -d '{ "clone": {
"source": { "name": "/vol/vol1/namespace2" } } }'
```

---

#### Converting a LUN into an NVMe namespace

An existing LUN can be converted in-place to an NVMe namespace with no modification to the data blocks. In other words, there is no additional copy created for the data blocks. There are certain requirements when converting a LUN to an NVMe namespace. For instance, the LUN should not be mapped to an initiator group, or exist as a protocol endpoint LUN, or in a foreign LUN import relationship. If the LUN exists as a VM volume, it should not be bound to a protocol endpoint LUN. Furthermore, only LUN with a supported operating system type for NVMe namespace can be converted.

The conversion process updates the metadata to the LUN, making it an NVMe namespace. The conversion is both time and space efficient. After conversion, the new namespace behaves as a regular namespace and may be mapped to an NVMe subsystem.

### Convert a LUN into an NVMe namespace

You convert a LUN into an NVMe namespace by calling a POST to /storage/namespaces. Set `convert.lun.uuid` or `convert.lun.name` to identify the source LUN which is to be converted in-place into an NVMe namespace.

```
# The API:
POST /api/storage/namespaces

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/namespaces' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "convert": { "lun":
{ "name": "/vol/vol1/lun1" } } }'
```

---

#### Deleting an NVMe namespace

```
# The API:
DELETE /api/storage/namespaces/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/storage/namespaces/5c254d22-96a6-
42ac-aad8-0cd9ebd126b6' -H 'Accept: application/hal+json'
```

---

### Retrieve NVMe namespaces

GET `/storage/namespaces`

**Introduced In:** 9.6

Retrieves NVMe namespaces.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `auto_delete`
- `space.physical_used`
- `space.physical_used_by_snapshots`
- `space.efficiency_ratio`
- `space.snapshot.*`
- `subsystem_map.*`
- `status.mapped`
- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver nvme namespace show`
- `vserver nvme subsystem map show`

#### Learn more

- DOC /storage/namespaces to learn more and examples.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| auto_delete | boolean | query | False | Filter by auto_delete |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| status.state | string | query | False | Filter by status.state |
| status.read_only | boolean | query | False | Filter by status.read_only |
| status.mapped | boolean | query | False | Filter by status.mapped |
| status.container_state | string | query | False | Filter by status.container_state |
| create_time | string | query | False | Filter by create_time<br>• Introduced in: 9.7 |
| subsystem_map.anagrpid | string | query | False | Filter by subsystem_map.anagrpid |
| subsystem_map.nsid | string | query | False | Filter by subsystem_map.nsid |
| subsystem_map.subsystem.os_type | string | query | False | Filter by subsystem_map.subsystem.os_type<br>• Introduced in: 9.16 |
| subsystem_map.subsystem.name | string | query | False | Filter by subsystem_map.subsystem.name<br>• maxLength: 64<br>• minLength: 1 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| subsystem_map.subsystem.comment | string | query | False | Filter by subsystem_map.subsystem.comment<br>• Introduced in: 9.16<br>• maxLength: 255<br>• minLength: 0 |
| subsystem_map.subsystem.uuid | string | query | False | Filter by subsystem_map.subsystem.uuid |
| subsystem_map.subsystem.hosts.tls.key_type | string | query | False | Filter by subsystem_map.subsystem.hosts.tls.key_type<br>• Introduced in: 9.16 |
| subsystem_map.subsystem.hosts.nqn | string | query | False | Filter by subsystem_map.subsystem.hosts.nqn<br>• Introduced in: 9.16 |
| subsystem_map.subsystem.hosts.dh_hmac_chap.group_size | string | query | False | Filter by subsystem_map.subsystem.hosts.dh_hmac_chap.group_size<br>• Introduced in: 9.16 |
| subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function | string | query | False | Filter by subsystem_map.subsystem.hosts.dh_hmac_chap.hash_function<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subsystem_map.subsystem.hosts.dh_hmac_chap.mode | string | query | False | Filter by subsystem_map.subsystem.hosts.dh_hmac_chap.mode<br><br>• Introduced in: 9.16 |
| subsystem_map.subsystem.hosts.priority | string | query | False | Filter by subsystem_map.subsystem.hosts.priority<br><br>• Introduced in: 9.16 |
| subsystem_map.subsystem.hosts.proximity.local_svm | boolean | query | False | Filter by subsystem_map.subsystem.hosts.proximity.local_svm<br><br>• Introduced in: 9.17 |
| subsystem_map.subsystem.hosts.proximity.peer_svms.name | string | query | False | Filter by subsystem_map.subsystem.hosts.proximity.peer_svms.name<br><br>• Introduced in: 9.17 |
| subsystem_map.subsystem.hosts.proximity.peer_svms.uuid | string | query | False | Filter by subsystem_map.subsystem.hosts.proximity.peer_svms.uuid<br><br>• Introduced in: 9.17 |
| os_type | string | query | False | Filter by os_type |
| consistency_group.name | string | query | False | Filter by consistency_group.name<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| consistency_group.uuid | string | query | False | Filter by consistency_group.uuid<br><br>• Introduced in: 9.16 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.8 |
| name | string | query | False | Filter by name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| enabled | boolean | query | False | Filter by enabled |
| location.volume.name | string | query | False | Filter by location.volume.name |
| location.volume.uuid | string | query | False | Filter by location.volume.uuid |
| location.qtree.name | string | query | False | Filter by location.qtree.name |
| location.qtree.id | integer | query | False | Filter by location.qtree.id<br>• Max value: 4994<br>• Min value: 0 |
| location.node.name | string | query | False | Filter by location.node.name<br>• Introduced in: 9.10 |
| location.node.uuid | string | query | False | Filter by location.node.uuid<br>• Introduced in: 9.10 |
| location.namespace | string | query | False | Filter by location.namespace |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br>• Introduced in: 9.8 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.8 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.8 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.8 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.8 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.8 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.8 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.8 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.8 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |
| comment | string | query | False | Filter by comment<br>• maxLength: 254<br>• minLength: 0 |
| uuid | string | query | False | Filter by uuid |
| space.used | integer | query | False | Filter by space.used |
| space.efficiency_ratio | number | query | False | Filter by space.efficiency_ratio<br>• Introduced in: 9.16 |
| space.physical_used_by_snapshots | integer | query | False | Filter by space.physical_used_by_snapshots<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| space.guarantee.requested | boolean | query | False | Filter by space.guarantee.requested |
| space.guarantee.reserved | boolean | query | False | Filter by space.guarantee.reserved |
| space.block_size | integer | query | False | Filter by space.block_size |
| space.physical_used | integer | query | False | Filter by space.physical_used<br>• Introduced in: 9.16 |
| space.size | integer | query | False | Filter by space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[nvme_namespace] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "consistency_group": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cg1",
        "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "location": {
        "namespace": "namespace1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "qtree": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
```

```
          },
          "id": 1,
          "name": "qt1"
        },
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "/vol/volume1/qtree1/namespace1",
      "os_type": "string",
      "space": {
        "block_size": 512,
        "efficiency_ratio": 2.5,
        "physical_used": 1073741824,
        "physical_used_by_snapshots": 1073741824,
```

```
        "size": 1073741824,
        "used": 0
      },
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "status": {
        "container_state": "string",
        "state": "online"
      },
      "subsystem_map": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "anagrpid": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "comment": "string",
          "hosts": [
            {
              "dh_hmac_chap": {
                "group_size": "string",
                "hash_function": "string",
```

```
                "mode": "bidirectional"
              },
              "nqn": "nqn.1992-01.example.com:string",
              "priority": "string",
              "proximity": {
                "peer_svms": [
                  {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "peer1",
                    "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                  }
                ]
              },
              "tls": {
                "key_type": "configured"
              }
            }
          ],
          "name": "subsystem1",
          "os_type": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

source

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`.

consistency_group

The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.
A namespace's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

lun

The source LUN for convert operation. This can be specified using property `convert.lun.uuid` or `convert.lun.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST. A convert request from LUN to NVMe namespace cannot be combined with setting any other namespace properties. All other properties of the converted NVMe namespace come from the source LUN.

convert

This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired.

node

The cluster node that hosts the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

qtree

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the NVMe namespace within the ONTAP cluster.
NVMe namespaces do not support rename, or movement between volumes. Valid in POST.

| Name | Type | Description |
|---|---|---|
| namespace | string | The base name component of the NVMe namespace. Valid in POST.<br><br>If properties `name` and `location.namespace` are specified in the same request, they must refer to the base name.<br><br>NVMe namespaces do not support rename. |
| node | node | The cluster node that hosts the NVMe namespace. |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | The qtree in which the NVMe namespace is optionally located. Valid in POST.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.<br><br>NVMe namespaces do not support rename. |
| volume | volume | The volume in which the NVMe namespace is located. Valid in POST.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>NVMe namespaces do not support movement between volumes. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

qos_policy

The QoS policy for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. If no `qos_policy` is provided, a policy is be set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the namespace based on the value specified. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

tiering

The tiering placement and policy definitions for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace in bytes. The default for namespaces with an `os_type` of *vmware* is *512*. All other namespaces default to *4096*.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |
| efficiency_ratio | number | The storage efficiency ratio of the namespace without snapshots. (Logical Used / Used)<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| physical_used | integer | The number of bytes consumed on the disk by the namespace, excluding snapshots.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the namespace's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not be made smaller using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Optional in POST. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.
These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.
See the `provisioning_options.auto` property for full details.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_namespace

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| create_time | string | The time the NVMe namespace was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| location | location | The location of the NVMe namespace within the ONTAP cluster.<br>NVMe namespaces do not support rename, or movement between volumes. Valid in POST.<br>• Introduced in: 9.6<br>• readCreate: 1 |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |

| Name | Type | Description |
| --- | --- | --- |
| space | space | The storage space related properties of the NVMe namespace. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | status | Status information about the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the NVMe namespace. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an NVMe namespace

POST `/storage/namespaces`

**Introduced In:** 9.6

Creates an NVMe namespace.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe namespace.
- `name`, `location.volume.name` or `location.volume.uuid` - Existing volume in which to create the NVMe namespace.
- `name` or `location.namespace` - Base name for the NVMe namespace.
- `os_type` - Operating system from which the NVMe namespace will be accessed. (Not used for clones, which are created based on the `os_type` of the source NVMe namespace.)
- `space.size` - Size for the NVMe namespace. (Not used for clones, which are created based on the size of the source NVMe namespace.)

#### Default property values

If not specified in POST, the following default property values are assigned:

- `auto_delete` - *false*
- `space.block_size` - *4096* ( *512* when 'os_type' is *vmware* )

#### Related ONTAP commands

- `volume file clone autodelete`
- `volume file clone create`

- `vserver nvme namespace convert-from-lun`
- `vserver nvme namespace create`

POST is asynchronous when creating a new namespace. It is synchronous when converting a LUN to a namespace via the `convert` property.

**Learn more**

- DOC /storage/namespaces

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br><br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.<br><br>When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| convert | convert | This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired. |
| create_time | string | The time the NVMe namespace was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| location | location | The location of the NVMe namespace within the ONTAP cluster.<br>NVMe namespaces do not support rename, or movement between volumes. Valid in POST.<br>• Introduced in: 9.6<br>• readCreate: 1 |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |

| Name | Type | Description |
| --- | --- | --- |
| space | space | The storage space related properties of the NVMe namespace. |
| status | status | Status information about the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the NVMe namespace. |

**Example request**


```
{
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "consistency_group": {
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "convert": {
    "lun": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "location": {
    "namespace": "namespace1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": {
      "id": 1,
      "name": "qt1"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "string",
  "provisioning_options": {
    "exclude_aggregates": [
      {
        "name": "aggr1"
      }
    ],
    "qos_policy": {
      "name": "performance",
```

```
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "snapshot_policy": {
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "storage_service": {
      "name": "string"
    },
    "tiering": {
      "control": "string",
      "object_stores": [
        {
          "name": "string"
        }
      ],
      "policy": "string"
    }
  },
  "space": {
    "block_size": 512,
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "subsystem_map": {
    "anagrpid": "00103050h",
    "nsid": "00000001h",
    "subsystem": {
      "comment": "string",
      "hosts": [
        {
          "dh_hmac_chap": {
            "controller_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "group_size": "string",
            "hash_function": "string",
            "host_secret_key": "DHHC-
1:00:ia6zGodOr4SEG0Zzaw398rpY0wqipUWj4jWjUh4HWUz6aQ2n:",
            "mode": "bidirectional"
```

```
          },
          "nqn": "nqn.1992-01.example.com:string",
          "priority": "string",
          "proximity": {
            "peer_svms": [
              {
                "name": "peer1",
                "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
              }
            ]
          },
          "tls": {
            "configured_psk": "NVMeTLSkey-
1:01:VRLbtnN9AQb2WXW3c9+wEf/DRLz0QuLdbYvEhwtdWwNf9LrZ:",
            "key_type": "configured"
          }
        }
      ],
      "name": "subsystem1",
      "os_type": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917927 | The specified volume was not found. |
| 918235 | A volume with the specified UUID was not found. |
| 918236 | The specified `location.volume.uuid` and `location.volume.name` do not refer to the same volume. |
| 1254197 | The LUN specified for conversion to a namespace is mapped. |
| 1260121 | Cloning a namespace to a volume different than the source volume is not supported. |
| 1260136 | The specified destination for a clone operation already exists as a LUN, namespace, or file. |
| 2621462 | The supplied SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |

| Error Code | Description |
| --- | --- |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5242927 | The specified qtree was not found. |
| 5242950 | The specified `location.qtree.id` and `location.qtree.name` do not refer to the same qtree. |
| 5374127 | The specified namespace name is invalid. |
| 5374140 | LUN has a non-zero prefix and/or suffix size. |
| 5374141 | LUN is part of a SnapMirror active sync relationship. |
| 5374156 | A protocol endpoint LUN cannot be converted to an NVMe namespace. |
| 5374157 | LUN in an SVM with MetroCluster configured cannot be converted to an NVMe namespace. |
| 5374158 | LUN contains an operating system type that is not supported for NVMe namespace. |
| 5374352 | An invalid name was provided for the NVMe namespace. |
| 5374858 | The volume specified by `name` is not the same as that specified by `location.volume`. |
| 5374860 | The qtree specified by `name` is not the same as that specified by `location.qtree`. |
| 5374861 | The NVME namespace base name specified by `name` is not the same as that specified by `location.name`. |
| 5374862 | No NVMe namespace path base name was provided for the namespace. |
| 5374876 | The LUN specified for conversion to a namespace was not found. |
| 5376461 | The specified namespace name is invalid. |
| 5376462 | The specified namespace name is too long. |
| 5376463 | The snapshot portion of the specified namespace name is too long. |
| 5440509 | No suitable storage can be found for the specified requirements. |
| 13565952 | The NVMe namespace clone request failed. |
| 72089636 | Creating NVMe namespaces with `os_type` AIX is not supported until the effective cluster version is 9.13.1 or later. |
| 72089720 | NVMe namespaces cannot be created in snapshots. |

| Error Code | Description |
| --- | --- |
| 72089721 | The volume specified is in a load sharing mirror relationship. Namespaces are not supported in load sharing mirrors. |
| 72089722 | A negative size was provided for the NVMe namespace. |
| 72089723 | The specified size is too small for the NVMe namespace. |
| 72089724 | The specified size is too large for the NVMe namespace. |
| 72089725 | A LUN or NVMe namespace already exists at the specified path. |
| 72089727 | NVMe namespaces cannot be created on an SVM root volume. |
| 72089728 | NVMe namespaces cannot be created on a FlexGroup volume. |
| 72089732 | An NVMe namespace name can only contain characters A-Z, a-z, 0-9, "-", ".", "_", "{" and "}". |
| 72090005 | The specified `clone.source.uuid` and `clone.source.name` do not refer to the same NVMe namespace. |
| 72090006 | The specified `clone.source` was not found. |
| 72090007 | The specified `clone.source` was not found. |
| 72090009 | An error occurred after successfully creating the NVMe namespace. Some properties were not set. |
| 72090012 | The property cannot be specified when creating an NVMe namespace clone. The `target` property of the error object identifies the property. |
| 72090013 | The property is required except when creating an NVMe namespace clone. The `target` property of the error object identifies the property. |
| 72090014 | No volume was specified for the NVMe namespace. |
| 72090015 | An error occurred after successfully creating the NVMe namespace preventing the retrieval of its properties. |
| 72090033 | The `clone.source.uuid` property is not supported when specifying a source NVMe namespace from a snapshot. |
| 72090039 | The property cannot be specified at the same time when creating an NVMe namespace as a clone. The `target` property of the error object identifies the other property given with clone. |

| Error Code | Description |
| --- | --- |
| 72090040 | The property cannot be specified when converting a LUN into an NVMe namespace. The `target` property of the error object identifies the property. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

source

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |
| uuid | string | The unique identifier of the clone source NVMe namespace.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`.

| Name | Type | Description |
| --- | --- | --- |
| source | source | The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.<br><br>Valid in POST to create a new NVMe namespace as a clone of the source.<br><br>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another. |

consistency_group

The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.
A namespace's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

lun

The source LUN for convert operation. This can be specified using property `convert.lun.uuid` or `convert.lun.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST. A convert request from LUN to NVMe namespace cannot be combined with setting any other namespace properties. All other properties of the converted NVMe namespace come from the source LUN.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the source LUN. Valid in POST.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional. |
| uuid | string | The unique identifier of the source LUN. Valid in POST. |

convert

This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired.

| Name | Type | Description |
| --- | --- | --- |
| lun | lun | The source LUN for convert operation. This can be specified using property `convert.lun.uuid` or `convert.lun.name`. If both properties are supplied, they must refer to the same LUN.<br><br>Valid in POST. A convert request from LUN to NVMe namespace cannot be combined with setting any other namespace properties. All other properties of the converted NVMe namespace come from the source LUN. |

node

The cluster node that hosts the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

qtree

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

| Name | Type | Description |
|---|---|---|
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the NVMe namespace within the ONTAP cluster.
NVMe namespaces do not support rename, or movement between volumes. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| namespace | string | The base name component of the NVMe namespace. Valid in POST.<br><br>If properties `name` and `location.namespace` are specified in the same request, they must refer to the base name.<br><br>NVMe namespaces do not support rename. |
| node | node | The cluster node that hosts the NVMe namespace. |
| qtree | qtree | The qtree in which the NVMe namespace is optionally located. Valid in POST.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.<br><br>NVMe namespaces do not support rename. |
| volume | volume | The volume in which the NVMe namespace is located. Valid in POST.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>NVMe namespaces do not support movement between volumes. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

self_link

qos_policy

The QoS policy for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. If no `qos_policy` is provided, a policy is be set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the namespace based on the value specified. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| auto | boolean | If the volume specified in the request does not exist, automatically provision one of appropriate size. If the volume does exist, resize it to accommodate the new namespace.<br><br>This property is only supported on Unified ONTAP.<br><br>The following behavior is different from a traditional POST request:<br><br>• The operation is asynchronous.<br>• The `qos_policy` property is supported and is applied to the provisioned volume. A default QoS policy is applied to the volume if one is not provided.<br>• The `provisioning_options.count` property is supported, provisioning *count* namespaces on the volume using the specified properties.<br>• The `subsystem_map` property is supported. If the specified subsystem does not exist, it is created. The namespace is mapped to this subsystem. If a subsystem is provisioned in this way, it is deleted after it is no longer mapped to any namespaces.<br>• The `clone` and `convert` properties are not supported.<br>• When performing `records` based operations, specifying this property in the query applies to the entire operation. Specifying it for an individual record within the request applies to only that record.<br>• Many other `provisioning_options` properties are supported to control the placement of the namespace and the properties of the volume |

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of namespaces to provision with these properties. Only POST requests based on `space.size` are supported. When provided, the name is considered a prefix, and a suffix of the form _<*N*> is generated where N is the next available numeric index, starting with 1. |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |
| qos_policy | qos_policy | The QoS policy for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. If no `qos_policy` is provided, a policy is be set based on the `storage_service.name`, which defaults to the most performant service available. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| storage_service | storage_service | Determines the placement of the namespace based on the value specified. Valid in POST. |
| tiering | tiering | The tiering placement and policy definitions for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. |

| Name | Type | Description |
| --- | --- | --- |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning the volume to host the namespace. Only mirrored aggregates are used if this parameter is set to *true* and only unmirrored aggregates are used if this parameter is set to *false*. The default value is *true* for a MetroCluster configuration and is *false* for a non-MetroCluster configuration. |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace in bytes. The default for namespaces with an `os_type` of *vmware* is *512*. All other namespaces default to *4096*.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |
| efficiency_ratio | number | The storage efficiency ratio of the namespace without snapshots. (Logical Used / Used)<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |

| Name | Type | Description |
| --- | --- | --- |
| physical_used | integer | The number of bytes consumed on the disk by the namespace, excluding snapshots.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the namespace's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not be made smaller using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| controller_secret_key | string | The controller secret for NVMe in-band authentication. The value of this property is used by the NVMe host to authenticate the NVMe controller while establishing a connection. If unset, the controller is not authenticated. When supplied, the property `host_secret_key` must also be supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a controller secret has been set for the host, but the controller secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| host_secret_key | string | The host secret for NVMe in-band authentication. The value of this property is used by the NVMe controller to authenticate the NVMe host while establishing a connection. If unset, no authentication is performed by the host or controller. This property must be supplied if any other NVMe in-band authentication properties are supplied. Optional in POST.<br><br>This property is write-only. The `mode` property can be used to identify if a host secret has been set for the host, but the host secret value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br><br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| configured_psk | string | A user supplied pre-shared key (PSK) value in PSK Interchange Format. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree. This property is only allowed when `key_type` is `configured`. If `configured_psk` is supplied and `key_type` is unset, `key_type` defaults to `configured`.<br><br>This property is write-only. The `key_type` property can be used to identify if a configured PSK has been set for the host, but the PSK value cannot be read. To change the value, the host must be deleted from the subsystem and re-added. |

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Optional in POST. |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.
These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.
See the `provisioning_options.auto` property for full details.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_namespace

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| clone | clone | This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.<br><br>When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`. |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| convert | convert | This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| location | location | The location of the NVMe namespace within the ONTAP cluster.<br>NVMe namespaces do not support rename, or movement between volumes. Valid in POST.<br>• Introduced in: 9.6<br>• readCreate: 1 |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| space | space | The storage space related properties of the NVMe namespace. |
| status | status | Status information about the NVMe namespace. |

| Name | Type | Description |
| --- | --- | --- |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the NVMe namespace. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve historical performance metrics for an NVMe namespace

GET `/storage/namespaces/{nvme_namespace.uuid}/metrics`

**Introduced In:** 9.14

Retrieves historical space and performance metrics for an NVMe namespace.

#### Related ONTAP commands

- `statistics namespace show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| timestamp | string | query | False | Filter by timestamp |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |
| uuid | string | query | False | Filter by uuid |
| nvme_namespace.uuid | string | path | True | The unique identifier of the NVMe namespace. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the NVMe namespace. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve historical performance metrics for an NVMe namespace for a specific time

GET `/storage/namespaces/{nvme_namespace.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical space and performance metrics for an NVMe namespace for a specific time.

#### Related ONTAP commands

- `statistics namespace show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nvme_namespace.uuid | string | path | True | The unique identifier of the NVMe namespace. |
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the NVMe namespace. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
|---|---|---|
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an NVMe namespace**

DELETE `/storage/namespaces/{uuid}`

**Introduced In:** 9.6

Deletes an NVMe namespace.

#### Related ONTAP commands

- `vserver nvme namespace delete`

#### Learn more

- DOC /storage/namespaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe namespace to delete. |
| allow_delete_while_mapped | boolean | query | False | Allows deletion of a mapped NVMe namespace.<br>A mapped NVMe namespace might be in use. Deleting a mapped namespace also deletes the namespace map and makes the data no longer available, possibly causing a disruption in the availability of data. **This parameter should be used with caution.**<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The specified namespace was not found. |
| 72089796 | The namespace must be unmapped before deletion. |
| 72090016 | The namespace's aggregate is offline. The aggregate must be online to modify or remove the namespace. |
| 72090017 | The namespace's volume is offline. The volume must be online to modify or remove the namespace. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an NVMe namespace

GET `/storage/namespaces/{uuid}`

**Introduced In:** 9.6

Retrieves an NVMe namespace.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `auto_delete`
- `space.physical_used`
- `space.physical_used_by_snapshots`
- `space.efficiency_ratio`
- `space.snapshot.*`
- `subsystem_map.*`
- `status.mapped`
- `statistics.*`

- `metric.*`

#### Related ONTAP commands

- `vserver nvme namespace show`
- `vserver nvme subsystem map show`

#### Learn more

- DOC /storage/namespaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe namespace to retrieve. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| create_time | string | The time the NVMe namespace was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| location | location | The location of the NVMe namespace within the ONTAP cluster.<br>NVMe namespaces do not support rename, or movement between volumes. Valid in POST.<br>• Introduced in: 9.6<br>• readCreate: 1 |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| os_type | string | The operating system type of the NVMe namespace.<br><br>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone. |

| Name | Type | Description |
| --- | --- | --- |
| space | space | The storage space related properties of the NVMe namespace. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | status | Status information about the NVMe namespace. |
| subsystem_map | subsystem_map | The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.<br><br>There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the NVMe namespace. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "consistency_group": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "location": {
    "namespace": "namespace1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 1,
      "name": "qt1"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
```

```
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "string",
  "space": {
    "block_size": 512,
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
    "size": 1073741824,
    "used": 0
  },
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
```

```
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "subsystem_map": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "anagrpid": "00103050h",
    "nsid": "00000001h",
    "subsystem": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "hosts": [
        {
          "dh_hmac_chap": {
            "group_size": "string",
            "hash_function": "string",
            "mode": "bidirectional"
          },
          "nqn": "nqn.1992-01.example.com:string",
          "priority": "string",
          "proximity": {
            "peer_svms": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
```

```
                  },
                  "name": "peer1",
                  "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
                }
              ]
            },
            "tls": {
              "key_type": "configured"
            }
          }
        ],
        "name": "subsystem1",
        "os_type": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 72090006 | The specified namespace was not found. |
| 72090007 | The specified namespace was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

source

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`.

consistency_group

The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.
A namespace's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

lun

The source LUN for convert operation. This can be specified using property `convert.lun.uuid` or `convert.lun.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST. A convert request from LUN to NVMe namespace cannot be combined with setting any other namespace properties. All other properties of the converted NVMe namespace come from the source LUN.

convert

This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired.

node

The cluster node that hosts the NVMe namespace.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

qtree

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the NVMe namespace within the ONTAP cluster.
NVMe namespaces do not support rename, or movement between volumes. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| namespace | string | The base name component of the NVMe namespace. Valid in POST.<br><br>If properties `name` and `location.namespace` are specified in the same request, they must refer to the base name.<br><br>NVMe namespaces do not support rename. |
| node | node | The cluster node that hosts the NVMe namespace. |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | The qtree in which the NVMe namespace is optionally located. Valid in POST.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.<br><br>NVMe namespaces do not support rename. |
| volume | volume | The volume in which the NVMe namespace is located. Valid in POST.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>NVMe namespaces do not support movement between volumes. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

qos_policy

The QoS policy for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. If no `qos_policy` is provided, a policy is be set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the namespace based on the value specified. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

tiering

The tiering placement and policy definitions for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
| --- | --- | --- |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| block_size | integer | The size of blocks in the namespace in bytes. The default for namespaces with an `os_type` of *vmware* is *512*. All other namespaces default to *4096*.<br><br>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.<br>Valid in POST. |
| efficiency_ratio | number | The storage efficiency ratio of the namespace without snapshots. (Logical Used / Used)<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| physical_used | integer | The number of bytes consumed on the disk by the namespace, excluding snapshots.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the namespace's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not be made smaller using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| group_size | string | The Diffie-Hellman group size for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `2048_bit`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| hash_function | string | The hash function for NVMe in-band authentication. When property `host_secret_key` is provided, this property defaults to `sha_256`. When supplied, the property `host_secret_key` must also be supplied. Optional in POST. |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

| Name | Type | Description |
| --- | --- | --- |
| key_type | string | The method by which the TLS pre-shared key (PSK) is configured for the host. Optional in POST.<br><br>The values for property `key_type` and property `configured_psk` must logically agree.<br><br>Possible values:<br><br>• `none` - TLS is not configured for the host connection. No value is allowed for property `configured_psk`.<br>• `configured` - A user supplied PSK is configured for the NVMe/TCP-TLS transport connection between the host and the NVMe subsystem. A valid value for property `configured_psk` is required.<br><br>This property defaults to `none` unless a value is supplied for `configured_psk` in which case it defaults to `configured`. |

hosts

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| dh_hmac_chap | consistency_group_nvme_host_dh_hmac_chap | A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host. |
| nqn | string | The NVMe qualified name (NQN) used to identify the NVMe storage target. |

| Name | Type | Description |
| --- | --- | --- |
| priority | string | The host priority setting allocates appropriate NVMe I/O queues (count and depth) for the host to submit I/O commands. Absence of this property in GET implies io_queue count and I/O queue depth are being used. |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |
| tls | tls | A container for the configuration for NVMe/TCP-TLS transport session for the host. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A configurable comment for the NVMe subsystem. Optional in POST. |
| hosts | array[hosts] | The NVMe hosts configured for access to the NVMe subsystem. Optional in POST. |
| name | string | The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the NVMe subsystem's hosts. Optional in POST. |
| uuid | string | The unique identifier of the NVMe subsystem. |

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.
These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.
See the `provisioning_options.auto` property for full details.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| subsystem | subsystem | The NVMe subsystem to which the NVMe namespace is mapped. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an NVMe namespace

PATCH `/storage/namespaces/{uuid}`

**Introduced In:** 9.6

Updates an NVMe namespace.

#### Related ONTAP commands

- `volume file clone autodelete`
- `vserver nvme namespace modify`

**Learn more**

- DOC /storage/namespaces

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the NVMe namespace to update. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| clone | clone | This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.<br><br>When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`. |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| create_time | string | The time the NVMe namespace was created. |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| space | space | The storage space related properties of the NVMe namespace. |

| Name | Type | Description |
|---|---|---|
| status | status | Status information about the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

**Example request**

```
{
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "consistency_group": {
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "name": "/vol/volume1/qtree1/namespace1",
  "space": {
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374127 | The specified namespace name is invalid. |
| 5376461 | The specified namespace name is invalid. |
| 5376462 | The specified namespace name is too long. |
| 5376463 | The snapshot portion of the specified namespace name is too long. |
| 5376466 | An attempt was made to rename an NVMe namespace to a snapshot name. |
| 5376467 | An attempt was made to rename a primary NVMe namespace to a secondary name. |
| 5376468 | An attempt was made to rename an NVMe namespace to a reserved name. |
| 13565952 | The namespace clone request failed. |
| 72089724 | The specified namespace size is too large. |
| 72089730 | The specified namespace cannot be updated as it resides in a snapshot. |
| 72090005 | The specified `clone.source.uuid` and `clone.source.name` do not refer to the same LUN. |
| 72090006 | The specified namespace was not found. This can apply to `clone.source` or the target namespace. The `target` property of the error object identifies the property. |
| 72090007 | The specified namespace was not found. This can apply to `clone.source` or the target namespace. The `target` property of the error object identifies the property. |
| 72090010 | An error occurred after successfully overwriting data for the namespace as a clone. Some properties were not modified. |
| 72090011 | An error occurred after successfully modifying some of the properties of the namespace. Some properties were not modified. |

| Error Code | Description |
| --- | --- |
| 72090016 | The namespace's aggregate is offline. The aggregate must be online to modify or remove the namespace. |
| 72090017 | The namespace's volume is offline. The volume must be online to modify or remove the namespace. |
| 72090038 | An attempt was made to reduce the size of the specified namespace. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

source

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |
| uuid | string | The unique identifier of the clone source NVMe namespace.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`.

| Name | Type | Description |
| --- | --- | --- |
| source | source | The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.<br><br>Valid in POST to create a new NVMe namespace as a clone of the source.<br><br>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another. |

consistency_group

The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.
A namespace's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

lun

The source LUN for convert operation. This can be specified using property `convert.lun.uuid` or `convert.lun.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST. A convert request from LUN to NVMe namespace cannot be combined with setting any other namespace properties. All other properties of the converted NVMe namespace come from the source LUN.

convert

This sub-object is used in POST to convert a valid in-place LUN to an NVMe namespace. Setting a property in this sub-object indicates that a conversion from the specified LUN to NVMe namespace is desired.

node

The cluster node that hosts the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

qtree

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the NVMe namespace within the ONTAP cluster.
NVMe namespaces do not support rename, or movement between volumes. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| node | node | The cluster node that hosts the NVMe namespace. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

self_link

qos_policy

The QoS policy for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume. If no `qos_policy` is provided, a policy is be set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the namespace based on the value specified. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for the volume provisioned to host the namespace. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

| Name | Type | Description |
|---|---|---|
| requested | boolean | The requested space reservation policy for the NVMe namespace. If *true*, a space reservation is requested for the namespace; if *false*, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.<br><br>The space reservation policy for an NVMe namespace is determined by ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| reserved | boolean | Reports if the NVMe namespace is space guaranteed.<br><br>This property is *true* if a space guarantee is requested and the containing volume and aggregate support the request. This property is *false* if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request. |

space

The storage space related properties of the NVMe namespace.

| Name | Type | Description |
|---|---|---|
| efficiency_ratio | number | The storage efficiency ratio of the namespace without snapshots. (Logical Used / Used)<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the NVMe namespace. |
| physical_used | integer | The number of bytes consumed on the disk by the namespace, excluding snapshots.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the namespace's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the namespace object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the NVMe namespace. Valid in POST and PATCH. The NVMe namespace size can be increased but not be made smaller using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the NVMe namespace.<br><br>This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the NVMe namespace.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available. |
| mapped | boolean | Reports if the NVMe namespace is mapped to an NVMe subsystem.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the NVMe namespace allows only read access. |
| state | string | The state of the NVMe namespace. Normal states for a namespace are *online* and *offline*. Other states indicate errors. |

consistency_group_nvme_host_dh_hmac_chap

A container for the configuration of NVMe in-band authentication using the DH-HMAC-CHAP protocol for a host.

| Name | Type | Description |
| --- | --- | --- |
| mode | string | The expected NVMe in-band authentication mode for the host. This property is an indication of which secrets are configured for the host. When set to:<br>• none: The host has neither the host nor controller secret configured, and no authentication is performed.<br>• unidirectional: The host has a host secret configured. The controller will authenticate the host.<br>• bidirectional: The host has both a host and controller secret configured. The controller will authenticate the host and the host will authenticate the controller. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.

These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the host is proximal to the SVM for which it is configured. |

| Name | Type | Description |
| --- | --- | --- |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the host is proximal. |

tls

A container for the configuration for NVMe/TCP-TLS transport session for the host.

hosts

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define the SVMs to which the host is proximal. This information is used to properly report active optimized and active non-optimized network paths using an NVMe controller. If no configuration has been specified for the host, the sub-object is not present in GET requests.<br><br>These properties apply to all instances of the host in the NVMe subsystem in the SVM and its peers. |

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added computational cost to retrieving property values for `subsystem_map`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.
These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.
See the `provisioning_options.auto` property for full details.

| Name | Type | Description |
| --- | --- | --- |
| anagrpid | string | The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.<br><br>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |
| nsid | string | The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.<br><br>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nvme_namespace

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem’s hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. An NVMe namespace can then be resized or cloned. An NVMe namespace

cannot be renamed, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace.

| Name | Type | Description |
|---|---|---|
| auto_delete | boolean | This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.<br><br>When set to *true*, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new NVMe namespace is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.<br><br>When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH: `auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The namespace's consistency group. This property is populated for namespaces that are members of a consistency group. If the namespace is a member of a child consistency group, the parent consistency group is reported.<br>A namespace's consistency group is the consistency group of its containing volume. |
| create_time | string | The time the NVMe namespace was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, check the `status.state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created. |
| name | string | The name of the NVMe namespace.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Renaming an NVMe namespace is not supported. Valid in POST. |
| space | space | The storage space related properties of the NVMe namespace. |
| status | status | Status information about the NVMe namespace. |
| uuid | string | The unique identifier of the NVMe namespace. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Name-services

## ONTAP REST API Name-services endpoints

### Overview

ONTAP uses name-services to obtain information about users and clients. This information is used to authenticate users accessing data on or administering the storage system, and to map user credentials in a mixed environment. If the user database is stored in NIS or LDAP servers, NIS and LDAP name services need to be configured in ONTAP.
DNS is used for resolving the hostnames. ns-switch is used to configure the SVMs with sources to search for network information and the order in which to search them.

## Manage group membership cache settings

### Manage group membership cache settings

#### Overview

This API is used to retrieve and manage group-membership cache settings.

#### Examples

### Retrieving group-membership cache settings

---

The following example shows how to use the cache group-membership settings GET endpoint to retrieve group-membership cache settings.

```
# The API:
GET /api/name-services/cache/group-membership/settings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/group-
membership/settings?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
      "name": "vs43"
    },
    "enabled": true,
    "ttl": "PT1H"
  },
  {
    "svm": {
      "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
      "name": "vs34"
    },
    "enabled": true,
    "ttl": "PT5M"
  }
],
"num_records": 2
}
```

#### Retrieving group-membership cache settings for a given SVM

The following example shows how to use the cache group-membership settings GET endpoint to retrieve group-membership cache settings for a given SVM.

```
# The API:
GET /api/name-services/cache/group-membership/settings/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/group-
membership/settings/dc458b2f-2035-11ec-bfe2-005056bb6bef" -H "accept:
application/hal+json"

# The response:
{
"svm": {
  "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
  "name": "vs34"
},
"enabled": true,
"ttl": "PT1H"
}
```

#### Updating a group-membership cache setting

The following example shows how to use the cache group-membership settings PATCH endpoint to update a group-membership setting.

```
# The API:
PATCH /api/name-services/cache/group-membership/settings/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/cache/unix-
user/settings/02c9e252-41be-11e9-81d5-00a0986138f9" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
\"false\" }"
```

#### Retrieve group membership cache settings

GET `/name-services/cache/group-membership/settings`

**Introduced In:** 9.11

Retrieves group-membership cache settings.

#### Related ONTAP commands

- `vserver services name-service cache group-membership settings show`

#### Learn more

- DOC /name-services/cache/group-membership/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| enabled | boolean | query | False | Filter by enabled |
| ttl | string | query | False | Filter by ttl |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[group_membership_settings] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "ttl": "PT2H30M"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_membership_settings

Group membership cache setting.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Specifies whether or not the cache is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
|---|---|---|
| ttl | string | Specifies Time to Live for list of group members, in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a group membership cache setting for an SVM

GET `/name-services/cache/group-membership/settings/{svm.uuid}`

**Introduced In:** 9.11

Retrieves a group-membership cache setting for a given SVM.

#### Related ONTAP commands

- `vserver services name-service cache group-membership settings show`

#### Learn more

- DOC /name-services/cache/group-membership/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| enabled | boolean | query | False | Filter by enabled |
| ttl | string | query | False | Filter by ttl |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Specifies whether or not the cache is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live for list of group members, in ISO 8601 format. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT2H30M"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a group membership cache setting

PATCH `/name-services/cache/group-membership/settings/{svm.uuid}`

**Introduced In:** 9.11

Updates a group-membership cache setting.

#### Important notes

- svm.uuid field cannot be empty.
- Returns success in case no values are provided for update.

#### Related ONTAP commands

- `vserver services name-service cache group-membership settings modify`

#### Learn more

- DOC /name-services/cache/group-membership/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether or not the cache is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live for list of group members, in ISO 8601 format. |

#### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT2H30M"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

group_membership_settings

Group membership cache setting.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether or not the cache is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live for list of group members, in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage host cache settings

### Manage host cache settings

#### Overview

This API is used to retrieve and manage hosts cache settings.

#### Examples

### Retrieving hosts cache settings

---

The following examples shows how to use the cache host settings GET endpoint to retrieve host cache settings.

```
# The API:
GET /api/name-services/cache/host/settings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-
services/cache/host/settings?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
    "svm": {
      "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
      "name": "vs43"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "dns_ttl_enabled": true
  },
  {
    "uuid": "951e8676-2035-11ec-bfe2-005056bb6bef",
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "dns_ttl_enabled": true
  },
  {
    "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
    "svm": {
      "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
      "name": "vs34"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "dns_ttl_enabled": true
  }
],
"num_records": 3
}
```

### Retrieving host cache settings for a given SVM

The following examples shows how to use the cache hosts settings GET endpoint to retrieve hosts cache settings for a specific SVM.

```
# The API:
GET /api/name-services/cache/host/settings/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-
services/cache/host/settings/dc458b2f-2035-11ec-bfe2-005056bb6bef" -H
"accept: application/hal+json"

# The response:
{
"uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
"svm": {
  "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
  "name": "vs34"
},
"enabled": false,
"negative_cache_enabled": true,
"ttl": "P1D",
"negative_ttl": "PT1M",
"dns_ttl_enabled": true
}
```

### Updating a hosts cache setting

The following example shows how to use the cache host settings PATCH endpoint to update a host cache setting.

```
# The API:
PATCH /api/name-services/cache/host/settings/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-
services/cache/host/settings/02c9e252-41be-11e9-81d5-00a0986138f9" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"ttl\": \"PT2H\", \"negative_ttl\": \"PT2M\" }"
```

### Retrieve host cache settings

GET `/name-services/cache/host/settings`

**Introduced In:** 9.11

Retrieves host cache settings.

#### Related ONTAP commands

- `vserver services name-service cache hosts settings show`

#### Learn more

- DOC /name-services/cache/host/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | query | False | Filter by uuid<br>• Introduced in: 9.14 |
| negative_ttl | string | query | False | Filter by negative_ttl |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| ttl | string | query | False | Filter by ttl |
| enabled | boolean | query | False | Filter by enabled |
| dns_ttl_enabled | boolean | query | False | Filter by dns_ttl_enabled |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[hosts_settings] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "dns_ttl_enabled": 1,
      "negative_ttl": "PT5M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "ttl": "PT24H",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

hosts_settings

Hosts cache setting.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| dns_ttl_enabled | boolean | Specifies whether TTL is based on the DNS or host settings. If enabled, TTL from DNS is used. |
| enabled | boolean | Indicates whether or not the cache is enabled. |

| Name | Type | Description |
|---|---|---|
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |
| uuid | string | UUID for the host record. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve a host cache setting for an SVM**

GET `/name-services/cache/host/settings/{uuid}`

**Introduced In:** 9.14

Retrieves a host cache setting for a given SVM.

#### Related ONTAP commands

- `vserver services name-service cache hosts settings show`

#### Learn more

- DOC /name-services/cache/host/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID for the host record. |
| negative_ttl | string | query | False | Filter by negative_ttl |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| ttl | string | query | False | Filter by ttl |
| enabled | boolean | query | False | Filter by enabled |
| dns_ttl_enabled | boolean | query | False | Filter by dns_ttl_enabled |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| dns_ttl_enabled | boolean | Specifies whether TTL is based on the DNS or host settings. If enabled, TTL from DNS is used. |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
|---|---|---|
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |
| uuid | string | UUID for the host record. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "dns_ttl_enabled": 1,
  "negative_ttl": "PT5M",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT24H",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a host cache setting

PATCH `/name-services/cache/host/settings/{uuid}`

**Introduced In:** 9.14

Updates a host cache setting.

#### Important notes

- UUID field cannot be empty.
- Returns success in case no values are provided for update.

#### Related ONTAP commands

- `vserver services name-service cache hosts settings modify`

#### Learn more

- DOC /name-services/cache/host/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID for the host record. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| dns_ttl_enabled | boolean | Specifies whether TTL is based on the DNS or host settings. If enabled, TTL from DNS is used. |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | UUID for the host record. |

#### Example request

```
{
  "dns_ttl_enabled": 1,
  "negative_ttl": "PT5M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT24H",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

hosts_settings

Hosts cache setting.

| Name | Type | Description |
|---|---|---|
| dns_ttl_enabled | boolean | Specifies whether TTL is based on the DNS or host settings. If enabled, TTL from DNS is used. |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |
| uuid | string | UUID for the host record. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage netgroup cache settings

### Manage netgroup cache settings

#### Overview

This API is used to retrieve and manage netgroups cache settings.

#### Examples

### Retrieving a netgroup cache setting

---

The following example shows how to use the cache netgroup settings GET endpoint to retrieve netgroup cache settings.

```
# The API:
GET /api/name-services/cache/netgroup/settings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-
services/cache/netgroup/settings?fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
      "name": "vs43"
    },
    "enabled": true,
    "negative_cache_enabled_byhost": true,
    "ttl_byhost": "P1D",
    "negative_ttl_byhost": "PT30M",
    "ttl_for_members": "PT30M"
  },
  {
    "svm": {
      "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
      "name": "vs34"
    },
    "enabled": true,
    "negative_cache_enabled_byhost": true,
    "ttl_byhost": "PT2H",
    "negative_ttl_byhost": "PT30M",
    "ttl_for_members": "PT30M"
  }
],
"num_records": 2
}
```

#### Retrieving a netgroup cache setting for a given SVM

The following example shows how to use the cache netgroup settings GET endpoint to retrieve netgroup cache settings for a given SVM.

```
# The API:
GET /api/name-services/cache/netgroup/settings/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-
services/cache/netgroup/settings/dc458b2f-2035-11ec-bfe2-005056bb6bef" -H
"accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
  "name": "vs34"
},
"enabled": true,
"negative_cache_enabled_byhost": true,
"ttl_byhost": "P1D",
"negative_ttl_byhost": "PT30M",
"ttl_for_members": "PT30M"
}
```

#### Updating a netgroup cache setting

The following example shows how to use the cache netgroup settings PATCH endpoint to update netgroup cache settings.

```
# The API:
PATCH /api/name-services/cache/netgroup/settings/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-
services/cache/netgroup/settings/02c9e252-41be-11e9-81d5-00a0986138f9" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"ttl_byhost\": \"PT2H\", \"negative_ttl_byhost\": \"PT2M\" }"
```

#### Retrieve netgroup cache settings

GET `/name-services/cache/netgroup/settings`

**Introduced In:** 9.11

Retrieves netgroups cache settings.

#### Related ONTAP commands

- `vserver services name-service cache netgroups settings show`

#### Learn more

- DOC /name-services/cache/netgroup/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ttl_for_members | string | query | False | Filter by ttl_for_members |
| enabled | boolean | query | False | Filter by enabled |
| negative_ttl_byhost | string | query | False | Filter by negative_ttl_byhost |
| negative_cache_enabled_byhost | boolean | query | False | Filter by negative_cache_enabled_byhost |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| ttl_byhost | string | query | False | Filter by ttl_byhost |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[netgroups_settings] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "negative_ttl_byhost": "PT2M5S",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "ttl_byhost": "PT24H",
      "ttl_for_members": "PT2M"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

netgroups_settings

Netgroups cache setting.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled_byhost | boolean | Indicates whether or not the negative cache is enabled. |

| Name | Type | Description |
|---|---|---|
| negative_ttl_byhost | string | Specifies negative Time to Live by host, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl_byhost | string | Specifies Time to Live by host, in ISO 8601 format. |
| ttl_for_members | string | Specifies Time to Live for netgroup members, in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a netgroup cache setting for an SVM

GET `/name-services/cache/netgroup/settings/{svm.uuid}`

**Introduced In:** 9.11

Retrieves a netgroup cache setting for a given SVM.

#### Related ONTAP commands

- `vserver services name-service cache netgroups settings show`

**Learn more**

- DOC /name-services/cache/netgroup/settings

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |
| ttl_for_members | string | query | False | Filter by ttl_for_members |
| enabled | boolean | query | False | Filter by enabled |
| negative_ttl_byhost | string | query | False | Filter by negative_ttl_byhost |
| negative_cache_enabled_byhost | boolean | query | False | Filter by negative_cache_enabled_byhost |
| svm.name | string | query | False | Filter by svm.name |
| ttl_byhost | string | query | False | Filter by ttl_byhost |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled_byhost | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl_byhost | string | Specifies negative Time to Live by host, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl_byhost | string | Specifies Time to Live by host, in ISO 8601 format. |

| Name | Type | Description |
| --- | --- | --- |
| ttl_for_members | string | Specifies Time to Live for netgroup members, in ISO 8601 format. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "negative_ttl_byhost": "PT2M5S",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl_byhost": "PT24H",
  "ttl_for_members": "PT2M"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a netgroup cache setting

PATCH `/name-services/cache/netgroup/settings/{svm.uuid}`

**Introduced In:** 9.11

Updates a netgroup cache setting.

#### Important notes

- svm.uuid field cannot be empty.
- Returns success in case no values are provided for update.

#### Related ONTAP commands

- `vserver services name-service cache netgroups settings modify`

#### Learn more

- DOC /name-services/cache/netgroup/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled_byhost | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl_byhost | string | Specifies negative Time to Live by host, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl_byhost | string | Specifies Time to Live by host, in ISO 8601 format. |
| ttl_for_members | string | Specifies Time to Live for netgroup members, in ISO 8601 format. |

#### Example request

```
{
  "negative_ttl_byhost": "PT2M5S",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl_byhost": "PT24H",
  "ttl_for_members": "PT2M"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

netgroups_settings

Netgroups cache setting.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled_byhost | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl_byhost | string | Specifies negative Time to Live by host, in ISO 8601 format. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl_byhost | string | Specifies Time to Live by host, in ISO 8601 format. |
| ttl_for_members | string | Specifies Time to Live for netgroup members, in ISO 8601 format. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage global cache settings

### Manage global cache settings

#### Overview

This API is used to retrieve and manage global nameservice cache settings.

#### Examples

### Retrieving a global nameservice cache setting

---

The following example shows how to use the cache setting GET endpoint to retrieve the global nameservice cache setting.

```
# The API:
GET /api/name-services/cache/setting

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/setting" -H
"accept: application/hal+json"

# The response:
{
"eviction_time_interval": "P2D",
"remote_fetch_enabled": true
}
```

---

### Updating a global nameservice cache setting

---

The following example shows how to use the cache setting PATCH endpoint to update the global nameservice cache setting.

```
# The API:
PATCH /api/name-services/cache/setting

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/cache/setting" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"eviction_time_interval\": \"PT2H\" }"
```

---

### Retrieve a global name service cache setting

GET `/name-services/cache/setting`

**Introduced In:** 9.11

Retrieves a global nameservice cache setting.

#### Related ONTAP commands

- `vserver services name-service cache settings show`

#### Learn more

- DOC /name-services/cache/setting

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| eviction_time_interval | string | query | False | Filter by eviction_time_interval |
| remote_fetch_enabled | boolean | query | False | Filter by remote_fetch_enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| eviction_time_interval | string | Specifies the time interval, in ISO 8601 format after which a periodic cache eviction happens. Default is 4 hours. |
| remote_fetch_enabled | boolean | Indicates whether or not a node is allowed to fetch the data from a remote node. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "eviction_time_interval": "PT2H"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a global name service cache setting

PATCH `/name-services/cache/setting`

**Introduced In:** 9.11

Updates a global nameservice cache setting.

#### Important notes

- Both the cache eviction time and remote fetch option can be modified.

#### Related ONTAP commands

- `vserver services name-service cache settings modify`

#### Learn more

- DOC /name-services/cache/setting

#### Request Body

| Name | Type | Description |
|---|---|---|
| eviction_time_interval | string | Specifies the time interval, in ISO 8601 format after which a periodic cache eviction happens. Default is 4 hours. |
| remote_fetch_enabled | boolean | Indicates whether or not a node is allowed to fetch the data from a remote node. |

#### Example request

```
{
  "eviction_time_interval": "PT2H"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

global_cache_setting

| Name | Type | Description |
|---|---|---|
| eviction_time_interval | string | Specifies the time interval, in ISO 8601 format after which a periodic cache eviction happens. Default is 4 hours. |
| remote_fetch_enabled | boolean | Indicates whether or not a node is allowed to fetch the data from a remote node. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage UNIX group cache settings

## Manage UNIX group cache settings

### Overview

This API is used to retrieve and manage unix-group settings.

### Examples

## Retrieving unix-group settings

---

The following example shows how to use the cache unix-group settings GET endpoint to retrieve unix-group cache settings.

```
# The API:
GET /api/name-services/cache/unix-group/settings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/unix-
group/settings?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
      "name": "vs43"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "propagation_enabled": true
  },
  {
    "svm": {
      "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
      "name": "vs34"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "propagation_enabled": true
  }
],
"num_records": 2
}
```

#### Retrieving a unix-group setting for a given SVM

The following example shows how to use the cache unix-group settings GET endpoint to retrieve unix-group cache settings for a given SVM.

```
# The API:
GET /api/name-services/cache/unix-group/settings/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/unix-
group/settings/dc458b2f-2035-11ec-bfe2-005056bb6bef" -H "accept:
application/hal+json"

# The response:
{
"svm": {
  "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
  "name": "vs34"
},
"enabled": true,
"negative_cache_enabled": true,
"ttl": "P1D",
"negative_ttl": "PT1M",
"propagation_enabled": true
}
```

#### Updating a unix-group setting

The following example shows how to use the cache unix-group settings PATCH endpoint to update unix-group cache settings for an SVM.

```
# The API:
PATCH /api/name-services/cache/unix-group/settings/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/cache/unix-
group/settings/02c9e252-41be-11e9-81d5-00a0986138f9" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"ttl\":
\"PT2H\", \"negative_ttl\": \"PT2M\" }"
```

#### Retrieve UNIX group cache settings

GET `/name-services/cache/unix-group/settings`

**Introduced In:** 9.11

Retrieves unix-group cache settings.

#### Related ONTAP commands

- `vserver services name-service cache unix-group settings show`

#### Learn more

- DOC /name-services/cache/unix-group/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| negative_ttl | string | query | False | Filter by negative_ttl |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| propagation_enabled | boolean | query | False | Filter by propagation_enabled |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| ttl | string | query | False | Filter by ttl |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[unix_group_settings] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "negative_ttl": "PT5M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "ttl": "PT24H"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_group_settings

UNIX groups cache setting.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |

| Name | Type | Description |
|---|---|---|
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a UNIX group cache setting for an SVM

GET `/name-services/cache/unix-group/settings/{svm.uuid}`

**Introduced In:** 9.11

Retrieves a unix-group cache setting for a given SVM.

#### Related ONTAP commands

- `vserver services name-service cache unix-group settings show`

**Learn more**

- DOC /name-services/cache/unix-group/settings

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |
| negative_ttl | string | query | False | Filter by negative_ttl |
| svm.name | string | query | False | Filter by svm.name |
| propagation_enabled | boolean | query | False | Filter by propagation_enabled |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| ttl | string | query | False | Filter by ttl |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "negative_ttl": "PT5M",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT24H"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update UNIX group cache settings

PATCH `/name-services/cache/unix-group/settings/{svm.uuid}`

**Introduced In:** 9.11

Updates unix-group cache settings.

#### Important notes

- svm.uuid field cannot be empty.
- Returns success in case no values are provided for update.

#### Related ONTAP commands

- `vserver services name-service cache unix-group settings modify`

#### Learn more

- DOC /name-services/cache/unix-group/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

#### Example request

```
{
  "negative_ttl": "PT5M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT24H"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_group_settings

UNIX groups cache setting.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage UNIX user cache settings

### Manage UNIX user cache settings

#### Overview

This API is used to retrieve and manage unix-user settings.

#### Examples

### Retrieving unix-user cache settings

---

The following example shows how to use the cache unix-user settings GET endpoint to retrieve unix-user cache settings.

```
# The API:
GET /api/name-services/cache/unix-user/settings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/unix-
user/settings?fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "8a1a8730-2036-11ec-8457-005056bbcfdb",
      "name": "vs43"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "propagation_enabled": true
  },
  {
    "svm": {
      "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
      "name": "vs34"
    },
    "enabled": true,
    "negative_cache_enabled": true,
    "ttl": "P1D",
    "negative_ttl": "PT1M",
    "propagation_enabled": true
  }
],
"num_records": 2
}
```

#### Retrieving a unix-user cache setting for a given SVM

The following example shows how to use the cache unix-user settings GET endpoint to retrieve unix-user cache settings for a given SVM.

```
# The API:
GET /api/name-services/cache/unix-user/settings/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/cache/unix-
user/settings/dc458b2f-2035-11ec-bfe2-005056bb6bef" -H "accept:
application/hal+json"

# The response:
{
"svm": {
  "uuid": "dc458b2f-2035-11ec-bfe2-005056bb6bef",
  "name": "vs34"
},
"enabled": true,
"negative_cache_enabled": true,
"ttl": "P1D",
"negative_ttl": "PT1M",
"propagation_enabled": true
}
```

#### Updating a unix-user setting

The following example shows how to use the cache unix-user settings PATCH endpoint to update unix-user cache settings.

```
# The API:
PATCH /api/name-services/cache/unix-user/settings/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/cache/unix-
user/settings/02c9e252-41be-11e9-81d5-00a0986138f9" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"ttl\":
\"PT2H\", \"negative_ttl\": \"PT2M\" }"
```

#### Retrieve UNIX user cache settings

GET `/name-services/cache/unix-user/settings`

**Introduced In:** 9.11

Retrieves unix-user cache settings.

#### Related ONTAP commands

- `vserver services name-service cache unix-user settings show`

#### Learn more

- DOC /name-services/cache/unix-user/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| negative_ttl | string | query | False | Filter by negative_ttl |
| propagation_enabled | boolean | query | False | Filter by propagation_enabled |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| enabled | boolean | query | False | Filter by enabled |
| ttl | string | query | False | Filter by ttl |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[unix_user_settings] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "negative_ttl": "PT5M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "ttl": "PT2H"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_user_settings

UNIX users cache setting.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |

| Name | Type | Description |
|---|---|---|
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve UNIX user cache settings for an SVM

GET `/name-services/cache/unix-user/settings/{svm.uuid}`

**Introduced In:** 9.11

Retrieves unix-user cache settings for a given SVM.

#### Related ONTAP commands

- `vserver services name-service cache unix-user settings show`

**Learn more**

- DOC /name-services/cache/unix-user/settings

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |
| svm.name | string | query | False | Filter by svm.name |
| negative_ttl | string | query | False | Filter by negative_ttl |
| propagation_enabled | boolean | query | False | Filter by propagation_enabled |
| negative_cache_enabled | boolean | query | False | Filter by negative_cache_enabled |
| enabled | boolean | query | False | Filter by enabled |
| ttl | string | query | False | Filter by ttl |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "negative_ttl": "PT5M",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT2H"
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a UNIX user cache setting

PATCH `/name-services/cache/unix-user/settings/{svm.uuid}`

**Introduced In:** 9.11

Updates a unix-user cache setting.

#### Important notes

- svm.uuid field cannot be empty.
- Returns success in case no values are provided for update.

#### Related ONTAP commands

- `vserver services name-service cache unix-user settings modify`

#### Learn more

- DOC /name-services/cache/unix-user/settings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

#### Example request

```
{
  "negative_ttl": "PT5M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ttl": "PT2H"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_user_settings

UNIX users cache setting.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the cache is enabled. |
| negative_cache_enabled | boolean | Indicates whether or not the negative cache is enabled. |
| negative_ttl | string | Specifies negative Time to Live, in ISO 8601 format. |
| propagation_enabled | boolean | Specifies whether the propagation setting is enabled or not. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| ttl | string | Specifies Time to Live (TTL), in ISO 8601 format. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage DNS configurations

### Manage DNS configurations

#### Overview

Displays DNS information and controls the DNS subsystem. DNS domain name and DNS servers are required parameters.

#### Retrieving DNS information

The DNS GET endpoint retrieves all of the DNS configurations for all SVMs.
DNS configuration for the cluster is retrieved via /api/cluster .

#### Examples

### Retrieving all of the fields for all of the DNS configurations

```
# The API:
/api/name-services/dns

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns?fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
```

```
  {
    "svm": {
      "uuid": "27eff5d8-22b2-11eb-8038-0050568ed32c",
      "name": "clust-1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domains": [
      "domain.example.com"
    ],
    "servers": [
      "44.44.44.44"
    ],
    "timeout":2,
    "attempts":1,
    "scope": "cluster",
    "_links": {
      "self": {
        "href": "/api/name-services/dns/27eff5d8-22b2-11eb-8038-
0050568ed32c"
      }
    }
  },
  {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domains": [
      "domainA.example.com"
    ],
    "servers": [
      "10.10.10.10"
    ],
    "timeout":2,
    "attempts":1,
    "dynamic_dns": {
```

```
      "enabled": false,
      "use_secure": false,
      "time_to_live": "PT1H"
    },
    "scope": "svm",
    "_links": {
      "self": {
        "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-
005056b41bd1"
      }
    }
  },
  {
    "uuid": "19076d35-6e27-11e8-b9b8-005056b41bd1",
    "svm": {
      "uuid": "19076d35-6e27-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19076d35-6e27-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domains": [
      "sample.example.com"
    ],
    "servers": [
      "11.11.11.11",
      "22.22.22.22",
      "33.33.33.33"
    ],
    "timeout":2,
    "attempts":2,
    "dynamic_dns": {
      "enabled": true,
      "use_secure": false,
      "time_to_live": "PT3H"
    },
    "scope": "svm",
    "_links": {
      "self": {
        "href": "/api/name-services/dns/19076d35-6e27-11e8-b9b8-
005056b41bd1"
      }
    }
  }
```

```
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/name-services/dns?fields=*"
  }
}
}
```


**Retrieving all DNS configurations whose domain name starts with *dom**.**

```
# The API:
/api/name-services/dns

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns?domains=dom*" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domains": [
      "domainA.example.com"
    ],
    "_links": {
      "self": {
        "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/name-services/dns?domains=dom*"
  }
}
}
```

**Retrieving the DNS configuration for a specific SVM**

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json"

# The response:
{
"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"domains": [
  "domainA.example.com"
],
"servers": [
  "10.10.10.10"
],
"timeout":2,
"attempts":1,
"tld_query_enabled": true,
"source_address_match": true,
"packet_query_match": true,
"dynamic_dns": {
  "enabled": false,
  "use_secure": false,
  "time_to_live": "P1D"
},
"scope": "svm",
"_links": {
  "self": {
    "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1"
  }
}
}
```

**Retrieving the advanced fields "DNS status", "tld_query_enable", "source_address_match", and "packet_query_match" using wildcards ****

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1?fileds=**" -H "accept: application/hal+json"

# The response:
{
"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"domains": [
  "domainA.example.com"
],
"servers": [
  "10.10.10.10"
],
"timeout":2,
"attempts":1,
"tld_query_enabled": true,
"source_address_match": true,
"packet_query_match": true,
"dynamic_dns": {
  "enabled": false,
  "use_secure": false,
  "time_to_live": "P1D"
},
"scope": "svm",
"status": [
  {
    "name_server": "10.10.10.10",
    "state": "up",
    "message": "Response time (msec): ",
    "code": 0
  }
```

```
],
"_links": {
  "self": {
    "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1"
  }
}
}
```

**Retrieving the "service_ips" for a specific DNS service name**

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-
005056b41bd1?service.name=_kpasswd._udp.domainA.example.com&fields=**" -H
"accept: application/hal+json"

# The response:
{
"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
"domains": [
  "domainA.example.com"
],
"servers": [
  "10.10.10.10"
],
"timeout": 2,
"attempts": 1,
"tld_query_enabled": true,
"source_address_match": true,
"packet_query_match": true,
"dynamic_dns": {
  "enabled": false,
  "use_secure": false,
  "time_to_live": "P1D"
},
"scope": "svm",
"status": [
  {
    "name_server": "10.10.10.10",
    "state": "up",
    "message": "Response time (msec): 218",
    "code": 0
```

```
    }
  ],
  "svm": {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
      }
    }
  },
  "service_ips": [
    "10.10.10.10",
    "2001:db08:a0b:12f0::1"
  ],
  "_links": {
    "self": {
      "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-
005056b41bd1?fields=**"
    }
  }
}
```

**Retrieving the "service_ips" for a specific DNS service name with a particular address type (can be ipv4, ipv6 or all)**

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-
005056b41bd1?service.name=_kpasswd._udp.domainA.example.com&service.addres
s_type=ipv4&fields=**" -H "accept: application/hal+json"

# The response:
{
"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
"domains": [
  "domainA.example.com"
],
"servers": [
  "10.10.10.10"
],
"timeout": 2,
```

```
"attempts": 1,
"tld_query_enabled": true,
"source_address_match": true,
"packet_query_match": true,
"dynamic_dns": {
  "enabled": false,
  "use_secure": false,
  "time_to_live": "P1D"
},
"scope": "svm",
"status": [
  {
    "name_server": "10.10.10.10",
    "state": "up",
    "message": "Response time (msec): 218",
    "code": 0
  }
],
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"service_ips": [
  "10.10.10.10"
],
"_links": {
  "self": {
    "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-
005056b41bd1?fields=**"
  }
}
}
```


#### Creating a DNS configuration

The DNS POST endpoint creates a DNS configuration for the specified SVM.

#### Examples

### Specifying only the required fields

```
# The API:
/api/name-services/dns

# The call:
    curl -X POST "https://<mgmt-ip>/api/name-services/dns" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domains\": [
\"domainA.example.com\" ], \"servers\": [ \"10.10.10.10\" ]}"
```

**Specifying the optional fields as well**

```
# The API:
/api/name-services/dns

# The call:
    curl -X POST "https://<mgmt-ip>/api/name-services/dns" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domains\": [
\"domainA.example.com\" ], \"servers\": [ \"10.10.10.10\" ],
\"timeout\":2, \"attempts\":3}"
```

**Specifying the scope of the SVM**

```
# The API:
/api/name-services/dns

# The call:
    curl -X POST "https://<mgmt-ip>/api/name-services/dns" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domains\": [
\"domainA.example.com\" ], \"servers\": [ \"10.10.10.10\" ],
\"timeout\":2, \"attempts\":3, \"scope\":\"svm\"}"
```

**Specifying the query parameter to run asynchronously.**

```
# The API:
/api/name-services/dns

# The call:
    curl -X POST "https://<mgmt-ip>/api/name-services/dns?async=true" -H
"accept: application/hal+json" -H "Content-Type: application/json" -d "{
\"svm\": { \"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" },
\"domains\": [ \"domainA.example.com\" ], \"servers\": [ \"10.10.10.10\"
]}"
```

#### Updating a DNS configuration

The DNS PATCH endpoint updates the DNS configuration for the specified SVM.

#### Examples

### Updating both the DNS domains and servers

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "{ \"domains\": [ \"domainA.example.com\",
\"domainB.example.com\" ], \"servers\": [ \"10.10.10.10\", \"10.10.10.11\"
]}"
```

### Updating the DNS servers only

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "{ \"servers\": [ \"10.10.10.10\" ]}"
```

### Updating the optional fields "timeout", "attempts", and "source_address_match"

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "{ \"timeout\":2, \"attempts\":3,
\"source_address_match\":true}"
```

**Updating the Dynamic DNS related fields**

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "{ \"timeout\":2, \"attempts\":3,
\"dynamic_dns.enabled\":\"true\", \"dynamic_dns.time_to_live\":\"20h\"}"
```

**Updating the DNS servers asynchronously.**

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1?async=true" -H "accept: application/hal+json" -H
"Content-Type: application/json" -d "{ \"servers\": [ \"10.10.10.10\" ]}"
```

#### Deleting a DNS configuration

The DNS DELETE endpoint deletes the DNS configuration for the specified SVM.

#### Example

The following example shows a DELETE operation.

```
# The API:
/api/name-services/dns/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-
11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"
```

### Retrieve DNS configurations for all SVMs

GET `/name-services/dns`

**Introduced In:** 9.6

Retrieves the DNS configurations of all SVMs.
Specify 'scope' as 'svm' to retrieve the DNS configuration of all the data SVMs.
Specify 'scope' as 'cluster' to retrieve the DNS configuration of the cluster.

#### Advanced properties

- 'tld_query_enabled'
- 'source_address_match'
- 'packet_query_match'
- 'status' property retrieves the status of each name server of the DNS configuration for an SVM.
- 'service_ips' property is displayed only when both service.name and SVM are set.

#### Related ONTAP commands

- `vserver services name-service dns show`
- `vserver services name-service dns check`
- `vserver services name-service dns dynamic-update show`
- `vserver services access-check dns srv-lookup`

#### Learn more

- DOC /name-services/dns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| servers | string | query | False | Filter by servers |
| status.message | string | query | False | Filter by status.message<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| status.name_server | string | query | False | Filter by status.name_server<br>• Introduced in: 9.9 |
| status.state | string | query | False | Filter by status.state<br>• Introduced in: 9.9 |
| status.code | integer | query | False | Filter by status.code<br>• Introduced in: 9.9 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| packet_query_match | boolean | query | False | Filter by packet_query_match<br>• Introduced in: 9.9 |
| service_ips | string | query | False | Filter by service_ips<br>• Introduced in: 9.14 |
| tld_query_enabled | boolean | query | False | Filter by tld_query_enabled<br>• Introduced in: 9.9 |
| timeout | integer | query | False | Filter by timeout<br>• Introduced in: 9.9<br>• Max value: 5<br>• Min value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scope | string | query | False | Filter by scope<br>• Introduced in: 9.9 |
| domains | string | query | False | Filter by domains<br>• maxLength: 255<br>• minLength: 1 |
| source_address_match | boolean | query | False | Filter by source_address_match<br>• Introduced in: 9.9 |
| uuid | string | query | False | Filter by uuid<br>• Introduced in: 9.13 |
| attempts | integer | query | False | Filter by attempts<br>• Introduced in: 9.9<br>• Max value: 4<br>• Min value: 1 |
| dynamic_dns.use_secure | boolean | query | False | Filter by dynamic_dns.use_secure<br>• Introduced in: 9.9 |
| dynamic_dns.enabled | boolean | query | False | Filter by dynamic_dns.enabled<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| dynamic_dns.time_to_live | string | query | False | Filter by dynamic_dns.time_to_live<br>• Introduced in: 9.9 |
| dynamic_dns.fqdn | string | query | False | Filter by dynamic_dns.fqdn<br>• Introduced in: 9.9 |
| service.name | string | query | False | The string required to look up the SRV record on the DNS server.<br>• Introduced in: 9.14 |
| service.address_type | string | query | False | The type of IP address to be looked up on the DNS server.<br>• Introduced in: 9.14<br>• Default value: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of DNS domain records. |
| records | array[dns] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "attempts": 1,
      "domains": [
        "example.com",
        "example2.example3.com"
      ],
      "dynamic_dns": {
        "fqdn": "example.com",
        "time_to_live": "P2D"
      },
      "scope": "string",
      "servers": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
      ],
      "service_ips": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
      ],
      "skip_config_validation": true,
      "status": [
        {
          "code": 6684732,
          "message": "string",
          "name_server": "10.10.10.10",
          "state": "up"
        }
      ],
```

```
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "timeout": 2,
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dynamic_dns

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable or disable Dynamic DNS (DDNS) updates for the specified SVM. |
| fqdn | string | Fully Qualified Domain Name (FQDN) to be used for dynamic DNS updates. |
| skip_fqdn_validation | boolean | Enable or disable FQDN validation. |
| time_to_live | string | Time to live value for the dynamic DNS updates, in an ISO-8601 duration formatted string. Maximum Time To Live is 720 hours(P30D in ISO-8601 format) and the default is 24 hours(P1D in ISO-8601 format). |
| use_secure | boolean | Enable or disable secure dynamic DNS updates for the specified SVM. |

status

Status of each of the DNS name server configured for the specified SVM.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Code corresponding to the status message. Code is 0 when the state is 'up'. |
| message | string | Detailed description of the validation state if the state is 'down' or<br>the response time of the DNS server if the state is 'up'. |
| name_server | string | The IP address of the DNS server. The address can be either an IPv4 or an IPv6 address. |
| state | string | The validation status of the DNS server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

dns

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| dynamic_dns | dynamic_dns | |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |
| scope | string | Set to "svm" for DNS owned by an SVM, otherwise set to "cluster". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |

| Name | Type | Description |
|---|---|---|
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create DNS domain and server configurations

POST `/name-services/dns`

**Introduced In:** 9.6

Creates DNS domain and server configurations for an SVM.

#### Important notes

- Each SVM can have only one DNS configuration.
- The domain name and the servers fields cannot be empty.
- IPv6 must be enabled if IPv6 family addresses are specified in the `servers` field.
- Configuring more than one DNS server is recommended to avoid a single point of failure.
- The DNS server specified using the `servers` field is validated during this operation.

The validation fails in the following scenarios:

1. The server is not a DNS server.
2. The server does not exist.
3. The server is unreachable.
    - The DNS server validation can be skipped by setting the property "skip_config_validation" to "true".
    - Scope of the SVM can be specified using the "scope" parameter. "svm" scope refers to data SVMs and "cluster" scope refers to clusters.

### The following parameters are optional:

- timeout
- attempts
- source_address_match
- packet_query_match
- tld_query_enabled
- skip_config_validation
- scope
- async

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| async | boolean | query | False | If set to true, ONTAP creates and verifies the DNS configuration in the background, returning a job to monitor the result. Otherwise, ONTAP waits until after the configuration has been finalized to return a response to the client.<br>• Introduced in: 9.18<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br><br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br><br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |

| Name | Type | Description |
| --- | --- | --- |
| scope | string | Set to "svm" for DNS owned by an SVM, otherwise set to "cluster". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified DNS configuration is disabled. |
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

### Example request

```
{
  "attempts": 1,
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "scope": "string",
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "service_ips": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "status": [
    {
      "code": 6684732,
      "message": "string",
      "name_server": "10.10.10.10",
      "state": "up"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timeout": 2,
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 8847360 | Only admin or data SVMs allowed |
| 8847361 | Exceeded the maximum number of domains allowed. Maximum of six domains only |
| 8847362 | Exceeded the maximum number of name servers allowed. Maximum of three name servers only |
| 8847392 | Domain name cannot be an IP address |
| 8847393 | Top level domain name is invalid |
| 8847399 | One or more of the specified DNS servers do not exist or cannot be reached |
| 8847394 | FQDN name violated the limitations |
| 8847403 | Scope specified is invalid for the specified SVM |
| 9240587 | FQDN name cannot be empty |

| Error Code | Description |
| --- | --- |
| 23724157 | Servers list cannot be empty |
| 9240588 | FQDN name is too long. Maximum supported length: 255 characters |
| 9240590 | FQDN name is reserved. Following names are reserved: "all", "local" and "localhost" |
| 9240607 | One of the FQDN labels is too long. Maximum supported length is 63 characters |
| 13434916 | The SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |
| 1377682 | IPv6 is not enabled in the cluster |
| 8847407 | The connection is refused for one or more of the specified DNS servers. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dynamic_dns

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable or disable Dynamic DNS (DDNS) updates for the specified SVM. |
| fqdn | string | Fully Qualified Domain Name (FQDN) to be used for dynamic DNS updates. |
| skip_fqdn_validation | boolean | Enable or disable FQDN validation. |
| time_to_live | string | Time to live value for the dynamic DNS updates, in an ISO-8601 duration formatted string. Maximum Time To Live is 720 hours(P30D in ISO-8601 format) and the default is 24 hours(P1D in ISO-8601 format). |
| use_secure | boolean | Enable or disable secure dynamic DNS updates for the specified SVM. |

status

Status of each of the DNS name server configured for the specified SVM.

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the status message. Code is 0 when the state is 'up'. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Detailed description of the validation state if the state is 'down' or<br>the response time of the DNS server if the state is 'up'. |
| name_server | string | The IP address of the DNS server. The address can be either an IPv4 or an IPv6 address. |
| state | string | The validation status of the DNS server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

dns

| Name | Type | Description |
| --- | --- | --- |
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |
| scope | string | Set to "svm" for DNS owned by an SVM, otherwise set to "cluster". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |

| Name | Type | Description |
| --- | --- | --- |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified DNS configuration is disabled. |
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the DNS domain configuration for an SVM

DELETE `/name-services/dns/{uuid}`

**Introduced In:** 9.13

Deletes DNS domain configuration of the specified SVM.

#### Related ONTAP commands

- `vserver services name-service dns delete`

#### Learn more

- DOC /name-services/dns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the DNS domain and server configuration for an SVM

GET `/name-services/dns/{uuid}`

**Introduced In:** 9.13

Retrieves DNS domain and server configuration of an SVM. By default, both DNS domains and servers are displayed.

#### Advanced properties

- 'tld_query_enabled'
- 'source_address_match'
- 'packet_query_match'
- 'status' property retrieves the status of each name server of the DNS configuration for an SVM.
- 'service_ips' property is displayed only when both service.name and SVM are set.

#### Related ONTAP commands

- `vserver services name-service dns show`
- `vserver services name-service dns check`
- `vserver services name-service dns dynamic-update show`
- `vserver services access-check dns srv-lookup`

#### Learn more

- DOC /name-services/dns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the DNS object. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| dynamic_dns | dynamic_dns | |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |
| scope | string | Set to "svm" for DNS owned by an SVM, otherwise set to "cluster". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |

| Name | Type | Description |
| --- | --- | --- |
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "attempts": 1,
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "dynamic_dns": {
    "fqdn": "example.com",
    "time_to_live": "P2D"
  },
  "scope": "string",
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "service_ips": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "skip_config_validation": true,
  "status": [
    {
      "code": 6684732,
      "message": "string",
      "name_server": "10.10.10.10",
      "state": "up"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timeout": 2,
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dynamic_dns

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable or disable Dynamic DNS (DDNS) updates for the specified SVM. |
| fqdn | string | Fully Qualified Domain Name (FQDN) to be used for dynamic DNS updates. |
| skip_fqdn_validation | boolean | Enable or disable FQDN validation. |
| time_to_live | string | Time to live value for the dynamic DNS updates, in an ISO-8601 duration formatted string. Maximum Time To Live is 720 hours(P30D in ISO-8601 format) and the default is 24 hours(P1D in ISO-8601 format). |
| use_secure | boolean | Enable or disable secure dynamic DNS updates for the specified SVM. |

status

Status of each of the DNS name server configured for the specified SVM.

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the status message. Code is 0 when the state is 'up'. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Detailed description of the validation state if the state is 'down' or<br>the response time of the DNS server if the state is 'up'. |
| name_server | string | The IP address of the DNS server. The address can be either an IPv4 or an IPv6 address. |
| state | string | The validation status of the DNS server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Update the DNS domain and server configuration for an SVM

PATCH `/name-services/dns/{uuid}`

**Introduced In:** 9.13

Updates DNS domain and server configurations of an SVM.

#### Important notes

- Both DNS domains and servers can be modified.
- The domains and servers fields cannot be empty.
- IPv6 must be enabled if IPv6 family addresses are specified for the `servers` field.
- The DNS server specified using the `servers` field is validated during this operation.

The validation fails in the following scenarios:

1. The server is not a DNS server.
2. The server does not exist.
3. The server is unreachable.
    - The DNS server validation can be skipped by setting the property "skip_config_validation" to "true".
    - Dynamic DNS configuration can be modified.
    - If both DNS and Dynamic DNS parameters are modified, DNS parameters are updated first followed by Dynamic DNS parameters.
      If updating Dynamic DNS fails, then the updated DNS configuration is not reverted.

#### The following parameters are optional:

- timeout
- attempts
- source_address_match
- packet_query_match
- tld_query_enabled
- skip_config_validation
- dynamic_dns.enabled
- dynamic_dns.use_secure
- dynamic_dns.time_to_live
- async

#### Related ONTAP commands

- `vserver services name-service dns modify`
- `vserver services name-service dns dynamic-update modify`

#### Learn more

- DOC /name-services/dns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the DNS object. |
| async | boolean | query | False | If set to true, ONTAP creates and verifies the DNS configuration in the background, returning a job to monitor the result. Otherwise, ONTAP waits until after the configuration has been finalized to return a response to the client.<br>• Introduced in: 9.18<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| dynamic_dns | dynamic_dns | |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified DNS configuration is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

#### Example request

```
{
  "attempts": 1,
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "dynamic_dns": {
    "fqdn": "example.com",
    "time_to_live": "P2D"
  },
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "service_ips": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "status": [
    {
      "code": 6684732,
      "message": "string",
      "name_server": "10.10.10.10",
      "state": "up"
    }
  ],
  "timeout": 2,
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

#### Response

```
Status: 200, Ok
```

#### Response

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8847360 | Only admin or data SVMs allowed |
| 8847361 | Exceeded the maximum number of domains allowed. Maximum of six domains only |
| 8847362 | Exceeded the maximum number of name servers allowed. Maximum of three name servers only |
| 8847376 | FQDN is mandatory if dynamic DNS update is being enabled. |
| 8847380 | Secure updates can be enabled only after a CIFS server or an Active Directory account has been created for the SVM. |
| 8847381 | A unique FQDN must be specified for each SVM. |
| 8847383 | The specified TTL exceeds the maximum supported value of 720 hours. |
| 8847392 | Domain name cannot be an IP address |
| 8847393 | Top level domain name is invalid |
| 8847394 | FQDN name violated the limitations |
| 8847399 | One or more of the specified DNS servers do not exist or cannot be reached |
| 8847404 | Dynamic DNS is applicable only for data SVMs |
| 8847405 | DNS parameters updated successfully; however the update of Dynamic DNS-related parameters has failed. |
| 9240587 | FQDN name cannot be empty |
| 9240588 | FQDN name is too long. Maximum supported length: 255 characters |
| 9240590 | FQDN name is reserved. Following names are reserved: "all", "local" and "localhost" |
| 9240607 | One of the FQDN labels is too long. Maximum supported length is 63 characters |
| 9240608 | FQDN does not contain a ".". At least one "." is mandatory for an FQDN. |
| 9240607 | A label of the FQDN is too long (73 characters). Maximum supported length for each label is 63 characters. |

| Error Code | Description |
| --- | --- |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

dynamic_dns

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable or disable Dynamic DNS (DDNS) updates for the specified SVM. |
| fqdn | string | Fully Qualified Domain Name (FQDN) to be used for dynamic DNS updates. |
| skip_fqdn_validation | boolean | Enable or disable FQDN validation. |
| time_to_live | string | Time to live value for the dynamic DNS updates, in an ISO-8601 duration formatted string. Maximum Time To Live is 720 hours(P30D in ISO-8601 format) and the default is 24 hours(P1D in ISO-8601 format). |
| use_secure | boolean | Enable or disable secure dynamic DNS updates for the specified SVM. |

status

Status of each of the DNS name server configured for the specified SVM.

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the status message. Code is 0 when the state is 'up'. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Detailed description of the validation state if the state is 'down' or<br>the response time of the DNS server if the state is 'up'. |
| name_server | string | The IP address of the DNS server. The address can be either an IPv4 or an IPv6 address. |
| state | string | The validation status of the DNS server. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

dns

| Name | Type | Description |
| --- | --- | --- |
| attempts | integer | Number of attempts allowed when querying the DNS name servers. |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| dynamic_dns | dynamic_dns | |
| packet_query_match | boolean | Indicates whether or not the query section of the reply packet is equal to that of the query packet. |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |
| service_ips | array[string] | List of IP addresses for a DNS service. Addresses can be IPv4, IPv6 or both. |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified DNS configuration is disabled. |

| Name | Type | Description |
|---|---|---|
| source_address_match | boolean | Indicates whether or not the DNS responses are from a different IP address to the IP address the request was sent to. |
| status | array[status] | Status of all the DNS name servers configured for the specified SVM. |
| timeout | integer | Timeout values for queries to the name servers, in seconds. |
| tld_query_enabled | boolean | Enable or disable top-level domain (TLD) queries. |
| uuid | string | UUID of the DNS object. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View host records

### View host records

#### Overview

Displays the IP address of the specified hostname and vice-versa.

#### Retrieving the host table entries

The host-record GET endpoint to retrieve the hostname for a given Ip address and vice-versa.

#### Examples

**Retrieving the hostname for a given IP address.**

```
# The API:
GET /api/name-services/host-record/{svm.uuid}/{host}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/host-record/77e23bd4-
a8fe-11eb-99e0-0050568e14ff/127.0.0.1" -H "accept: application/json"

# The response:
 {
"svm": {
  "uuid": "77e23bd4-a8fe-11eb-99e0-0050568e14ff",
  "name": "svm1"
 },
"host": "127.0.0.1",
"ipv4_addresses": [
  "127.0.0.1"
 ],
"hostname": "localhost",
"source": "Files"
 }
```

**Retrieving the Ip address for a given hostname.**

```
# The API:
GET /api/name-services/host-record/{svm.uuid}/{host}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/host-record/77e23bd4-
a8fe-11eb-99e0-0050568e14ff/localhost" -H "accept: application/json"

# The response:
 {
"svm": {
  "uuid": "77e23bd4-a8fe-11eb-99e0-0050568e14ff",
  "name": "svm1"
},
"host": "localhost",
"ipv4_addresses": [
  "127.0.0.1"
],
"ipv6_addresses": [
  "::1"
],
"hostname": "localhost",
"canonical_name": "localhost",
"source": "Files"
 }
```

---

### Retrieve an IP address of a hostname

GET `/name-services/host-record/{svm.uuid}/{host}`

**Introduced In:** 9.10

Retrieves the IP address of the specified hostname and vice-versa.

#### Related ONTAP commands

- `vserver services name-service getxxbyyy getnameinfo`
- `vserver services name-service getxxbyyy getaddrinfo`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| host | string | path | True | Hostname or IP address. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| use_cache | boolean | query | False | Enables or disables the cache.<br>• Introduced in: 9.11<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| canonical_name | string | Canonical name of the host. |
| host | string | IP address or hostname. |
| hostname | string | Hostname. |
| ipv4_addresses | array[string] | List of IPv4 addresses. |
| ipv6_addresses | array[string] | List of IPv6 addresses. |
| source | string | Source used for lookup. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "canonical_name": "localhost",
  "host": "localhost",
  "hostname": "localhost",
  "ipv4_addresses": [
    "127.0.0.1"
  ],
  "ipv6_addresses": [
    ":1"
  ],
  "source": "Files",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8912900 | Invalid IP address. |
| 23724050 | Failed to resolve the specified hostname or IP address. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage LDAP server configurations

## Manage LDAP server configurations

### Overview

LDAP servers are used to centrally maintain user information. LDAP configurations must be set up to lookup information stored in the LDAP directory on the external LDAP servers. This API is used to retrieve and manage
LDAP server configurations.

### Retrieving LDAP information

The LDAP GET endpoint retrieves all of the LDAP configurations in the cluster.

### Examples

#### Retrieving all of the fields for all LDAP configurations

```
# The API:
/api/name-services/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap?fields=**" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "servers": [
      "10.10.10.10",
      "domainB.example.com"
    ],
    "schema": "ad_idmu",
    "port": 389,
    "ldaps_enabled": false,
    "min_bind_level": "anonymous",
```

```
    "bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
    "base_dn": "dc=domainA,dc=example,dc=com",
    "base_scope": "subtree",
    "use_start_tls": true,
    "session_security": "none",
    "referral_enabled": false,
    "bind_as_cifs_server": false,
    "query_timeout": 3,
    "is_owner": true,
    "user_scope": "subtree",
    "group_scope": "subtree",
    "netgroup_scope": "subtree",
    "is_netgroup_byhost_enabled": false,
    "netgroup_byhost_scope": "subtree",
    "group_membership_filter": "",
    "status": {
      "state": "down",
      "message": "The LDAP configuration is invalid. Verify that the AD
domain or servers are reachable and that the network configuration is
correct",
      "dn_message": [
        "No LDAP DN configured"
      ],
      "code": 4915258,
      "ipv4_state": "down",
      "ipv6_state": "down"
    },
    "try_channel_binding": true,
    "restrict_discovery_to_site" : false,
    "_links": {
      "self": {
        "href": "/api/name-services/ldap/179d3c85-7053-11e8-b9b8-
005056b41bd1"
      }
    }
  },
  {
    "svm": {
      "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    },
```

```
    "ad_domain": "example.com",
    "schema": "rfc_2307",
    "port": 389,
    "ldaps_enabled": false,
    "min_bind_level": "simple",
    "bind_dn": "cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
    "base_dn": "dc=domainB,dc=example,dc=com",
    "base_scope": "subtree",
    "use_start_tls": true,
    "session_security": "sign",
    "referral_enabled": false,
    "bind_as_cifs_server": false,
    "query_timeout": 0,
    "is_owner": true,
    "user_scope": "subtree",
    "group_scope": "subtree",
    "netgroup_scope": "subtree",
    "is_netgroup_byhost_enabled": false,
    "netgroup_byhost_scope": "subtree",
    "group_membership_filter": "",
    "status": {
      "state": "up",
      "message": "Successfully connected to LDAP server \"172.20.192.44\".
Successfully connected to LDAP server
\"fd20:8b1e:b255:5056:999:d9:516c:bf69\".",
      "dn_message": [
        "All the configured DNs are available."
      ],
      "code": 0,
      "ipv4_state": "up",
      "ipv6_state": "up"
    },
    "try_channel_binding": true,
    "restrict_discovery_to_site": true,
    "_links": {
      "self": {
        "href": "/api/name-services/ldap/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/name-services/ldap?fields=*"
```

```
    }
  }
  }
```

### Retrieving all of the LDAP configurations that have the *use_start_tls* set to *true*

```
# The API:
/api/name-services/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap?use_start_tls=true"
-H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "use_start_tls": true,
    "_links": {
      "self": {
        "href": "/api/name-services/ldap/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/name-services/ldap?use_start_tls=true"
  }
}
}
```

#### Retrieving the LDAP configuration of a specific SVM

```
# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"servers": [
  "10.10.10.10",
  "domainB.example.com"
],
"schema": "ad_idmu",
"port": 389,
"ldaps_enabled": false,
"min_bind_level": "anonymous",
"bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
"base_dn": "dc=domainA,dc=example,dc=com",
"base_scope": "subtree",
"use_start_tls": true,
"referral_enabled": false,
"session_security": "none",
"bind_as_cifs_server": true,
"query_timeout": 3,
"is_owner": true,
"try_channel_binding": true,
"restrict_discovery_to_site": false,
"_links": {
  "self": {
    "href": "/api/name-services/ldap/179d3c85-7053-11e8-b9b8-005056b41bd1"
  }
}
}
```

#### Retrieving all the fields of the LDAP configuration of a specific SVM

```
# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-11e8-
b9b8-005056b41bd1?fields=**" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"servers": [
  "10.10.10.10",
  "domainB.example.com"
],
"schema": "ad_idmu",
"port": 389,
"ldaps_enabled": false,
"min_bind_level": "anonymous",
"bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
"base_dn": "dc=domainA,dc=example,dc=com",
"base_scope": "subtree",
"use_start_tls": true,
"referral_enabled": false,
"session_security": "none",
"bind_as_cifs_server": true,
"query_timeout": 3,
"is_owner": true,
"user_scope": "subtree",
"group_scope": "subtree",
"netgroup_scope": "subtree",
"is_netgroup_byhost_enabled": false,
"netgroup_byhost_scope": "subtree",
"group_membership_filter": "",
"try_channel_binding": true,
```

```
"restrict_discovery_to_site": false,
"status": {
  "state": "down",
  "message": "The LDAP configuration is invalid. Verify that the AD domain
or servers are reachable and that the network configuration is correct",
  "dn_message": [
    "No LDAP DN configured"
  ],
  "code": 4915258,
  "ipv4_state": "down",
  "ipv6_state": "down"
},
"_links": {
  "self": {
    "href": "/api/name-services/ldap/179d3c85-7053-11e8-b9b8-005056b41bd1"
  }
}
}
```


### Retrieving the LDAP server status of a specific SVM

```
# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap/9e4a2e3b-f66f-11ea-
aec8-0050568e155c?fields=status" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "9e4a2e3b-f66f-11ea-aec8-0050568e155c",
  "name": "vs2"
},
"status": {
  "state": "up",
  "message": "Successfully connected to LDAP server \"172.20.192.44\".
Successfully connected to LDAP server
\"fd20:8b1e:b255:5056:999:d9:516c:bf69\".",
  "code": 0,
  "ipv4_state": "up",
  "ipv6_state": "up"
}
}
```

**Retrieving all of the LDAP configurations that have the *restrict_discovery_to_site* set to *true***

```
# The API:
/api/name-services/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/name-
services/ldap?restrict_discovery_to_site=true" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "restrict_discovery_to_site": true,
    "_links": {
      "self": {
        "href": "/api/name-services/ldap/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/name-services/ldap?restrict_discovery_to_site=true"
  }
}
}
```

#### Creating an LDAP configuration

The LDAP POST endpoint creates an LDAP configuration for the specified SVM.

#### Examples

### Creating an LDAP configuration with all the fields specified

```
# The API:
/api/name-services/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"servers\": [
\"10.10.10.10\", \"domainB.example.com\" ], \"schema\": \"ad_idmu\",
\"port\": 389, \"ldaps_enabled\": false, \"min_bind_level\":
\"anonymous\", \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"base_scope\": \"subtree\", \"use_start_tls\": false,
\"session_security\": \"none\", \"referral_enabled\": false,
\"bind_as_cifs_server\": false, \"query_timeout\": 4, \"user_dn\":
\"cn=abc,users,dc=com\", \"user_scope\": \"subtree\", \"group_dn\":
\"cn=abc,users,dc=com\", \"group_scope\": \"subtree\", \"netgroup_dn\":
\"cn=abc,users,dc=com\", \"netgroup_scope\": \"subtree\",
\"netgroup_byhost_dn\": \"cn=abc,users,dc=com\",
\"netgroup_byhost_scope\": \"subtree\", \"is_netgroup_byhost_enabled\":
false, \"group_membership_filter\": \"\", \"skip_config_validation\":
false } "
```

### Creating an LDAP configuration with Active Directory domain and preferred Active Directory servers specified

```
# The API:
/api/name-services/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"name\": \"vs2\" }, \"ad_domain\": \"domainA.example.com\",
\"preferred_ad_servers\": [ \"11.11.11.11\" ], \"port\": 389,
\"ldaps_enabled\": false, \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"session_security\": \"none\", \"referral_enabled\": false,
\"query_timeout\": 3}"
```

#### Creating an LDAP configuration with a number of optional fields not specified

```
# The API:
/api/name-services/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"name\": \"vs2\" }, \"servers\": [ \"11.11.11.11\" ], \"port\": 389,
\"bind_dn\": \"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"session_security\": \"none\"}"
```

#### Creating an LDAP configuration with Active Directory domain specified and *restrict_discovery_to_site* set to *true*

```
# The API:
/api/name-services/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"name\": \"vs2\" }, \"ad_domain\": \"example.com\", \"port\": 389,
\"bind_dn\": \"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"session_security\": \"none\", \"restrict_discovery_to_site\": true}"
```

#### Updating an LDAP configuration

The LDAP PATCH endpoint updates the LDAP configuration for the specified SVM. The following example shows a PATCH operation:

```
# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-
11e8-b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"servers\": [ \"55.55.55.55\" ], \"schema\":
\"ad_idmu\", \"port\": 636, \"ldaps_enabled\": true, \"use_start_tls\":
false, \"referral_enabled\": false ,\"restrict_discovery_to_site\":
false}"
```

#### Deleting an LDAP configuration

The LDAP DELETE endpoint deletes the LDAP configuration for the specified SVM. The following example shows a DELETE operation:

```
# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-
11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"
```

### Retrieve the LDAP configuration for all SVMs

GET `/name-services/ldap`

**Introduced In:** 9.6

Retrieves the LDAP configurations for all SVMs.

#### Related ONTAP commands

- `ldap show`
- `ldap check -vserver vs0`
- `ldap check-ipv6 -vserver vs0`

#### Important notes

- The status.code, status.dn_message, status.message, and status.state fields have the same status fields that are returned using the "ldap check" CLI command.
- Refer to the ipv4 or ipv6 objects available in the status field to get specific information about the code, dn_messages, or message and state information for ipv4 or ipv6.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| min_bind_level | string | query | False | Filter by min_bind_level<br>• Introduced in: 9.7 |
| bind_dn | string | query | False | Filter by bind_dn<br>• Introduced in: 9.7 |
| netgroup_dn | string | query | False | Filter by netgroup_dn<br>• Introduced in: 9.9 |
| bind_as_cifs_server | boolean | query | False | Filter by bind_as_cifs_server<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| restrict_discovery_to_site | boolean | query | False | Filter by restrict_discovery_to_site<br>• Introduced in: 9.13 |
| user_dn | string | query | False | Filter by user_dn<br>• Introduced in: 9.9 |
| use_start_tls | boolean | query | False | Filter by use_start_tls<br>• Introduced in: 9.7 |
| ldaps_enabled | boolean | query | False | Filter by ldaps_enabled<br>• Introduced in: 9.9 |
| netgroup_byhost_scope | string | query | False | Filter by netgroup_byhost_scope<br>• Introduced in: 9.9 |
| is_owner | boolean | query | False | Filter by is_owner<br>• Introduced in: 9.9 |
| base_scope | string | query | False | Filter by base_scope<br>• Introduced in: 9.7 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| svm.uuid | string | query | False | Filter by svm.uuid<br>• Introduced in: 9.7 |
| ad_domain | string | query | False | Filter by ad_domain<br>• Introduced in: 9.7 |
| netgroup_byhost_dn | string | query | False | Filter by netgroup_byhost_dn<br>• Introduced in: 9.9 |
| session_security | string | query | False | Filter by session_security<br>• Introduced in: 9.7 |
| user_scope | string | query | False | Filter by user_scope<br>• Introduced in: 9.9 |
| schema | string | query | False | Filter by schema<br>• Introduced in: 9.7 |
| base_dn | string | query | False | Filter by base_dn<br>• Introduced in: 9.7 |
| group_dn | string | query | False | Filter by group_dn<br>• Introduced in: 9.9 |
| group_scope | string | query | False | Filter by group_scope<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| netgroup_scope | string | query | False | Filter by netgroup_scope<br>• Introduced in: 9.9 |
| group_membership_filter | string | query | False | Filter by group_membership_filter<br>• Introduced in: 9.9 |
| port | integer | query | False | Filter by port<br>• Introduced in: 9.7<br>• Max value: 65535<br>• Min value: 1 |
| preferred_ad_servers | string | query | False | Filter by preferred_ad_servers<br>• Introduced in: 9.7 |
| try_channel_binding | boolean | query | False | Filter by try_channel_binding<br>• Introduced in: 9.10 |
| servers | string | query | False | Filter by servers<br>• Introduced in: 9.7 |
| status.ipv4_state | string | query | False | Filter by status.ipv4_state<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| status.ipv6.message | string | query | False | Filter by status.ipv6.message<br><br>• Introduced in: 9.14 |
| status.ipv6.dn_messages | string | query | False | Filter by status.ipv6.dn_messages<br><br>• Introduced in: 9.14 |
| status.ipv6.code | integer | query | False | Filter by status.ipv6.code<br><br>• Introduced in: 9.14 |
| status.ipv6.state | string | query | False | Filter by status.ipv6.state<br><br>• Introduced in: 9.14 |
| status.ipv4.message | string | query | False | Filter by status.ipv4.message<br><br>• Introduced in: 9.14 |
| status.ipv4.dn_messages | string | query | False | Filter by status.ipv4.dn_messages<br><br>• Introduced in: 9.14 |
| status.ipv4.code | integer | query | False | Filter by status.ipv4.code<br><br>• Introduced in: 9.14 |
| status.ipv4.state | string | query | False | Filter by status.ipv4.state<br><br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| status.state | string | query | False | Filter by status.state<br>• Introduced in: 9.9 |
| status.ipv6_state | string | query | False | Filter by status.ipv6_state<br>• Introduced in: 9.13 |
| status.message | string | query | False | Filter by status.message<br>• Introduced in: 9.9 |
| status.code | integer | query | False | Filter by status.code<br>• Introduced in: 9.9 |
| status.dn_message | string | query | False | Filter by status.dn_message<br>• Introduced in: 9.9 |
| referral_enabled | boolean | query | False | Filter by referral_enabled<br>• Introduced in: 9.9 |
| is_netgroup_byhost_enabled | boolean | query | False | Filter by is_netgroup_byhost_enabled<br>• Introduced in: 9.9 |
| query_timeout | integer | query | False | Filter by query_timeout<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of LDAP records. |
| records | array[ldap_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_domain": "example.com",
      "base_dn": "dc=domainB,dc=example,dc=com",
      "base_scope": "string",
      "bind_dn":
"cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
      "group_dn": "cn=abc,users,dc=com",
      "group_membership_filter": "",
      "group_scope": "string",
      "min_bind_level": "string",
      "netgroup_byhost_dn": "cn=abc,users,dc=com",
      "netgroup_byhost_scope": "string",
      "netgroup_dn": "cn=abc,users,dc=com",
      "netgroup_scope": "string",
      "port": 389,
      "preferred_ad_servers": [
        "11.11.11.11"
      ],
      "schema": "ad_idmu",
      "servers": [
        [
          "10.10.10.10",
          "domainB.example.com"
        ]
      ],
      "session_security": "string",
      "skip_config_validation": true,
      "status": {
```

```
        "code": 65537300,
        "dn_message": [
          "string"
        ],
        "ipv4": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv4_state": "string",
        "ipv6": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv6_state": "string",
        "message": "string",
        "state": "string"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "user_dn": "cn=abc,users,dc=com",
      "user_scope": "string"
    }
  ]
}
```


**Error**

```
Status: Default
```

Following error codes can be thrown as part of LDAP status information, if LDAP status is needed to be retrieved.

| Error Code | Description |
|---|---|
| 4915229 | DNS resolution failed due to an internal error. Contact technical support if this issue persists |
| 4915231 | DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured |
| 23724132 | DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured |
| 4915258 | The LDAP configuration is invalid. Verify that the Active Directory domain or servers are reachable and that the network configuration is correct |
| 4915263 | Failed to check the current status of LDAP server. Reason:<Reason for="" the="" failure=""></Reason> |
| 4915234 | The specified LDAP server or preferred Active Directory server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast |
| 4915265 | The specified bind password or bind DN is invalid |
| 4915264 | Certificate verification failed. Verify that a valid certificate is installed |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |
| dn_message | array[string] | |
| ipv4 | ipv4 | |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ldap_service

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |

| Name | Type | Description |
| --- | --- | --- |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| status | status | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
|---|---|---|
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the LDAP configuration for an SVM

POST `/name-services/ldap`

**Introduced In:** 9.6

Creates an LDAP configuration for an SVM.

#### Important notes

- Each SVM can have one LDAP configuration.
- The LDAP servers and Active Directory domain are mutually exclusive fields. These fields cannot be empty. At any point in time, either the LDAP servers or Active Directory domain must be populated.

- LDAP configuration with Active Directory domain cannot be created on an admin SVM.
- IPv6 must be enabled if IPv6 family addresses are specified.

**The following parameters are optional:**

- preferred AD servers
- schema
- port
- ldaps_enabled
- min_bind_level
- bind_password
- base_scope
- use_start_tls
- session_security
- referral_enabled
- bind_as_cifs_server
- query_timeout
- user_dn
- user_scope
- group_dn
- group_scope
- netgroup_dn
- netgroup_scope
- netgroup_byhost_dn
- netgroup_byhost_scope
- is_netgroup_byhost_enabled
- group_membership_filter
- skip_config_validation
- try_channel_binding
- restrict_discovery_to_site

Configuring more than one LDAP server is recommended to avoid a single point of failure.
Both FQDNs and IP addresses are supported for the "servers" field.
The Active Directory domain or LDAP servers are validated as part of this operation.

LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server or Active Directory domain is invalid.
3. The server or Active Directory domain is unreachable.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| ad_domain | string | This parameter specifies the name of the Active Directory domain used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |

| Name | Type | Description |
| --- | --- | --- |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
| --- | --- | --- |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example request**

```
{
  "ad_domain": "example.com",
  "base_dn": "dc=domainB,dc=example,dc=com",
  "base_scope": "string",
  "bind_dn": "cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
  "bind_password": "abc",
  "group_dn": "cn=abc,users,dc=com",
  "group_membership_filter": "",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "cn=abc,users,dc=com",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "cn=abc,users,dc=com",
  "netgroup_scope": "string",
  "port": 389,
  "preferred_ad_servers": [
    "11.11.11.11"
  ],
  "schema": "ad_idmu",
  "servers": [
    [
      "10.10.10.10",
      "domainB.example.com"
    ]
  ],
  "session_security": "string",
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv4_state": "string",
    "ipv6": {
      "code": 65537300,
      "dn_messages": [
        "string"
```

```
      ],
      "message": "string",
      "state": "string"
    },
    "ipv6_state": "string",
    "message": "string",
    "state": "string"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user_dn": "cn=abc,users,dc=com",
  "user_scope": "string"
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of LDAP records. |
| records | array[ldap_service] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "ad_domain": "example.com",
      "base_dn": "dc=domainB,dc=example,dc=com",
      "base_scope": "string",
      "bind_dn":
"cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
      "bind_password": "abc",
      "group_dn": "cn=abc,users,dc=com",
      "group_membership_filter": "",
      "group_scope": "string",
      "min_bind_level": "string",
      "netgroup_byhost_dn": "cn=abc,users,dc=com",
      "netgroup_byhost_scope": "string",
      "netgroup_dn": "cn=abc,users,dc=com",
      "netgroup_scope": "string",
      "port": 389,
      "preferred_ad_servers": [
        "11.11.11.11"
      ],
      "schema": "ad_idmu",
      "servers": [
        [
          "10.10.10.10",
          "domainB.example.com"
        ]
      ],
      "session_security": "string",
      "status": {
        "code": 65537300,
        "dn_message": [
          "string"
        ],
        "ipv4": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv4_state": "string",
```

```
        "ipv6": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv6_state": "string",
        "message": "string",
        "state": "string"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "user_dn": "cn=abc,users,dc=com",
      "user_scope": "string"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262186 | LDAP Servers cannot be used with Active Directory domain and/or preferred Active Directory servers |
| 2621488 | Invalid SVM context |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 4915203 | The specified LDAP schema does not exist |
| 262222 | The specified LDAP servers or preferred Active Directory servers contain duplicate server entries |

| Error Code | Description |
| --- | --- |
| 4915229 | DNS resolution failed due to an internal error. Contact technical support if this issue persists |
| 4915231 | DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured |
| 23724132 | DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured |
| 4915234 | The specified LDAP server or preferred Active Directory server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast |
| 4915248 | LDAP servers cannot be empty or "-". Specified Active Directory domain is invalid because it is empty or "-" or it contains either the special characters or "-" at the start or end of the domain) |
| 4915251 | STARTTLS and LDAPS cannot be used together |
| 4915257 | The LDAP configuration is invalid. Verify that bind-dn and bind password are correct |
| 4915258 | The LDAP configuration is invalid. Verify that the Active Directory domain or servers are reachable and that the network configuration is correct |
| 4915259 | LDAP configurations with Active Directory domains are not supported on admin SVM. |
| 4915265 | The specified bind password or bind DN is invalid |
| 4915264 | Certificate verification failed. Verify that a valid certificate is installed |
| 13434916 | The SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |
| 4915252 | LDAP Referral is not supported with STARTTLS, with session security levels sign, seal or with LDAPS. |
| 4915266 | LDAP site discovery restriction cannot be applied to a mixed version cluster. |
| 656484 | The default site needs to be specified to enable site restriction. |
| 4915206 | CIFS server is not configured for the vserver. LDAP client configuration requires CIFS server for binding. |
| 4915261 | Cannot use port "389" when "ldaps_enabled" is "true". |
| 4915255 | Base DN specified in the LDAP client configuration is not available. |

| Error Code | Description |
| --- | --- |
| 4915268 | The bind_as_cifs_server field cannot be set to true when the CIFS server is in workgroup mode. |
| 4915269 | The bind_as_cifs_server field cannot be set to true when the CIFS server is in realm mode. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |
| dn_message | array[string] | |
| ipv4 | ipv4 | |

| Name | Type | Description |
| --- | --- | --- |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ldap_service

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

| Name | Type | Description |
| --- | --- | --- |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the LDAP configuration for an SVM

DELETE `/name-services/ldap/{svm.uuid}`

**Introduced In:** 9.6

Deletes the LDAP configuration of the specified SVM. LDAP can be removed as a source from the ns-switch if LDAP is not used as a source for lookups.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the LDAP configuration for an SVM

GET `/name-services/ldap/{svm.uuid}`

**Introduced In:** 9.6

Retrieves LDAP configuration for an SVM. All parameters for the LDAP configuration are displayed by default.

#### Related ONTAP commands

- `ldap show`
- `ldap check -vserver vs0`
- `ldap check-ipv6 -vserver vs0`

#### Important notes

- The status.code, status.dn_message, status.message, and status.state fields have the same status fields that are returned using the "ldap check" CLI command.
- Refer to the ipv4 or ipv6 objects available in the status field to get specific information about the code, dn_messages, or message and state information for ipv4 or ipv6.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ad_domain | string | This parameter specifies the name of the Active Directory domain used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |
| base_dn | string | Specifies the default base DN for all searches. |

| Name | Type | Description |
| --- | --- | --- |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| status | status | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |

| Name | Type | Description |
| --- | --- | --- |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": "example.com",
  "base_dn": "dc=domainB,dc=example,dc=com",
  "base_scope": "string",
  "bind_dn": "cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
  "group_dn": "cn=abc,users,dc=com",
  "group_membership_filter": "",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "cn=abc,users,dc=com",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "cn=abc,users,dc=com",
  "netgroup_scope": "string",
  "port": 389,
  "preferred_ad_servers": [
    "11.11.11.11"
  ],
  "schema": "ad_idmu",
  "servers": [
    [
      "10.10.10.10",
      "domainB.example.com"
    ]
  ],
  "session_security": "string",
  "skip_config_validation": true,
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
```

```
      "ipv4_state": "string",
      "ipv6": {
        "code": 65537300,
        "dn_messages": [
          "string"
        ],
        "message": "string",
        "state": "string"
      },
      "ipv6_state": "string",
      "message": "string",
      "state": "string"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "user_dn": "cn=abc,users,dc=com",
    "user_scope": "string"
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |

| Name | Type | Description |
| --- | --- | --- |
| dn_message | array[string] | |
| ipv4 | ipv4 | |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the LDAP configuration for an SVM

PATCH `/name-services/ldap/{svm.uuid}`

**Introduced In:** 9.6

Updates an LDAP configuration of an SVM.

#### Important notes

- Both mandatory and optional parameters of the LDAP configuration can be updated.
- The LDAP servers and Active Directory domain are mutually exclusive fields. These fields cannot be empty. At any point in time, either the LDAP servers or Active Directory domain must be populated.
- IPv6 must be enabled if IPv6 family addresses are specified.

Configuring more than one LDAP server is recommended to avoid a single point of failure.
Both FQDNs and IP addresses are supported for the "servers" field.
The Active Directory domain or LDAP servers are validated as part of this operation.

LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server or Active Directory domain is invalid.
3. The server or Active Directory domain is unreachable

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |

| Name | Type | Description |
| --- | --- | --- |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example request**

```
{
  "ad_domain": "example.com",
  "base_dn": "dc=domainB,dc=example,dc=com",
  "base_scope": "string",
  "bind_dn": "cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
  "bind_password": "abc",
  "group_dn": "cn=abc,users,dc=com",
  "group_membership_filter": "",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "cn=abc,users,dc=com",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "cn=abc,users,dc=com",
  "netgroup_scope": "string",
  "port": 389,
  "preferred_ad_servers": [
    "11.11.11.11"
  ],
  "schema": "ad_idmu",
  "servers": [
    [
      "10.10.10.10",
      "domainB.example.com"
    ]
  ],
  "session_security": "string",
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv4_state": "string",
    "ipv6": {
      "code": 65537300,
      "dn_messages": [
        "string"
```

```
        ],
        "message": "string",
        "state": "string"
      },
      "ipv6_state": "string",
      "message": "string",
      "state": "string"
    },
    "user_dn": "cn=abc,users,dc=com",
    "user_scope": "string"
  }
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262186 | LDAP Servers cannot be used with Active Directory domain and/or preferred Active Directory servers |
| 2621488 | Invalid SVM context |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 4915203 | The specified LDAP schema does not exist |
| 262222 | The specified LDAP servers or preferred Active Directory servers contain duplicate server entries |
| 4915229 | DNS resolution failed due to an internal error. Contact technical support if this issue persists |
| 4915231 | DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured |
| 23724132 | DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured |

| Error Code | Description |
| --- | --- |
| 4915234 | The specified LDAP server or preferred Active Directory server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast |
| 4915248 | LDAP servers cannot be empty or "-". Specified Active Directory domain is invalid because it is empty or "-" or it contains either the special characters or "-" at the start or end of the domain. |
| 4915251 | STARTTLS and LDAPS cannot be used together |
| 4915257 | The LDAP configuration is invalid. Verify that the distinguished names and bind password are correct |
| 4915258 | The LDAP configuration is invalid. Verify that the Active Directory domain or servers are reachable and that the network configuration is correct |
| 4915259 | LDAP configurations with Active Directory domains are not supported on admin SVM. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |
| 4915252 | LDAP Referral is not supported with STARTTLS, with session security levels sign, seal or with LDAPS. |
| 4915266 | LDAP site discovery restriction cannot be applied to a mixed version cluster. |
| 656484 | The default site needs to be specified to enable site restriction. |
| 4915206 | CIFS server is not configured for the vserver. LDAP client configuration requires CIFS server for binding. |
| 4915244 | RPC failure occurred during validation of the LDAP configuration. |
| 4915268 | The bind_as_cifs_server field cannot be set to true when the CIFS server is in workgroup mode. |
| 4915269 | The bind_as_cifs_server field cannot be set to true when the CIFS server is in realm mode. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |
| dn_message | array[string] | |
| ipv4 | ipv4 | |

| Name | Type | Description |
| --- | --- | --- |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ldap_service

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

| Name | Type | Description |
| --- | --- | --- |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage LDAP schemas

### Manage LDAP schemas

#### Overview

ONTAP provides four default LDAP schemas. These are:

- MS-AD-BIS
  Based on RFC-2307bis. This is the preferred LDAP schema for most standard Windows 2012 and later LDAP deployments.
- AD-IDMU
  Based on Active Directory Identity Management for UNIX. This schema is appropriate for most Windows 2008, Windows 2012, and later AD servers.
- AD-SFU
  Based on Active Directory Services for UNIX. This schema is appropriate for most Windows 2003 and earlier AD servers.
- RFC-2307
  Based on RFC-2307 (an approach that uses LDAP as a network information service). This schema is appropriate for most UNIX AD servers.

#### Examples

### Retrieving LDAP schema information

```
# The API:
GET /api/name-services/ldap-schemas

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap-schemas"

# The response:
```

```
{
"records": [
{
  "owner": {
    "uuid": "eda950c6-0a0c-11ec-bfcf-0050568e9150",
    "name": "athiraacluster-1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/eda950c6-0a0c-11ec-bfcf-0050568e9150"
      }
    }
  },
  "name": "AD-IDMU",
  "_links": {
    "self": {
      "href": "/api/name-services/ldap-schemas/eda950c6-0a0c-11ec-bfcf-
0050568e9150/AD-IDMU"
    }
  }
},
{
  "owner": {
    "uuid": "eda950c6-0a0c-11ec-bfcf-0050568e9150",
    "name": "athiraacluster-1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/eda950c6-0a0c-11ec-bfcf-0050568e9150"
      }
    }
  },
  "name": "AD-SFU",
  "_links": {
    "self": {
      "href": "/api/name-services/ldap-schemas/eda950c6-0a0c-11ec-bfcf-
0050568e9150/AD-SFU"
    }
  }
},
{
  "owner": {
    "uuid": "eda950c6-0a0c-11ec-bfcf-0050568e9150",
    "name": "athiraacluster-1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/eda950c6-0a0c-11ec-bfcf-0050568e9150"
      }
```

```
    }
  },
  "name": "MS-AD-BIS",
  "_links": {
    "self": {
      "href": "/api/name-services/ldap-schemas/eda950c6-0a0c-11ec-bfcf-
0050568e9150/MS-AD-BIS"
    }
  }
},
{
  "owner": {
    "uuid": "eda950c6-0a0c-11ec-bfcf-0050568e9150",
    "name": "athiraacluster-1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/eda950c6-0a0c-11ec-bfcf-0050568e9150"
      }
    }
  },
  "name": "RFC-2307",
  "_links": {
    "self": {
      "href": "/api/name-services/ldap-schemas/eda950c6-0a0c-11ec-bfcf-
0050568e9150/RFC-2307"
    }
  }
}
],
"num_records": 4,
"_links": {
"self": {
  "href": "/api/name-services/ldap-schemas"
}
}
}
```


#### Retrieving LDAP schema information for a given SVM and "name"

```
# The API:
GET /api/name-services/ldap-schemas
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap-schemas/eda950c6-
0a0c-11ec-bfcf-0050568e9150/RFC-2307"

# The response:
{
"owner": {
"uuid": "eda950c6-0a0c-11ec-bfcf-0050568e9150",
"name": "svm1",
"_links": {
  "self": {
    "href": "/api/svm/svms/eda950c6-0a0c-11ec-bfcf-0050568e9150"
  }
}
},
"name": "RFC-2307",
"comment": "Schema based on RFC 2307 (read-only)",
"name_mapping": {
"account": {
  "unix": "unixAccount",
  "windows": "windowsAccount"
},
"windows_to_unix": {
  "attribute": "windowsAccount",
  "no_domain_prefix": false,
  "object_class": "posixAccount"
}
},
"rfc2307": {
"attribute": {
  "gecos": "gecos",
  "gid_number": "gidNumber",
  "home_directory": "homeDirectory",
  "login_shell": "loginShell",
  "uid": "uid",
  "uid_number": "uidNumber",
  "user_password": "userPassword"
},
"cn": {
  "group": "cn",
  "netgroup": "cn"
},
"member": {
  "nis_netgroup": "memberNisNetgroup",
  "uid": "memberUid"
},
```

```
"nis": {
  "netgroup_triple": "nisNetgroupTriple",
  "netgroup": "nisNetgroup",
  "object": "nisObject",
  "mapname": "nisMapName",
  "mapentry": "nisMapEntry"
},
"posix": {
  "account": "posixAccount",
  "group": "posixGroup"
}
},
"rfc2307bis": {
"enabled": false,
"group_of_unique_names": "groupOfUniqueNames",
"maximum_groups": 256,
"unique_member": "uniqueMember"
},
"scope": "cluster",
"global_schema": true,
"_links": {
"self": {
  "href": "/api/name-services/ldap-schemas/eda950c6-0a0c-11ec-bfcf-
0050568e9150/RFC-2307"
}
}
}
```


### Creating an LDAP schema

```
# The API:
POST /api/name-services/ldap-schemas

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap-schemas" -H
"accept: application/json" -H "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{ \"name\": \"schema\",\"template\":
{ \"name\": \"AD-IDMU\" }, \"owner\": { \"uuid\": \"52ba8197-0a23-11ec-
9622-0050568e9150\", \"name\" : \"svm1\" }}"
```

### Updating an LDAP schema

```
# The API:
PATCH /api/name-services/ldap-schemas/{owner.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/ldap-schemas/52ba8197-
0a23-11ec-9622-0050568e9150/schema" -H "authorization: Basic
<CREDENTIALS>" -H "Content-Type: application/json" -d "{ \"comment\":
\"This is a comment for schema\"}"
```

### Deleting an LDAP schema

```
# The API:
DELETE /api/name-services/ldap-schemas/{owner.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/ldap-schemas/52ba8197-
0a23-11ec-9622-0050568e9150/schema"
```

### Retrieve all LDAP schemas

GET `/name-services/ldap-schemas`

**Introduced In:** 9.11

Retrieves all the LDAP schemas.

#### Related ONTAP commands

- `vserver services name-service ldap client schema show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| global_schema | boolean | query | False | Filter by global_schema |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• maxLength: 32<br>• minLength: 1 |
| rfc2307bis.maximum_groups | integer | query | False | Filter by rfc2307bis.maximum_groups<br>• Max value: 1024<br>• Min value: 1 |
| rfc2307bis.enabled | boolean | query | False | Filter by rfc2307bis.enabled |
| rfc2307bis.group_of_unique_names | string | query | False | Filter by rfc2307bis.group_of_unique_names |
| rfc2307bis.unique_member | string | query | False | Filter by rfc2307bis.unique_member |
| rfc2307.attribute.home_directory | string | query | False | Filter by rfc2307.attribute.home_directory |
| rfc2307.attribute.user_password | string | query | False | Filter by rfc2307.attribute.user_password |
| rfc2307.attribute.gid_number | string | query | False | Filter by rfc2307.attribute.gid_number |
| rfc2307.attribute.login_shell | string | query | False | Filter by rfc2307.attribute.login_shell |
| rfc2307.attribute.gecos | string | query | False | Filter by rfc2307.attribute.gecos |
| rfc2307.attribute.uid | string | query | False | Filter by rfc2307.attribute.uid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rfc2307.attribute.uid_number | string | query | False | Filter by rfc2307.attribute.uid_number |
| rfc2307.posix.account | string | query | False | Filter by rfc2307.posix.account |
| rfc2307.posix.group | string | query | False | Filter by rfc2307.posix.group |
| rfc2307.member.uid | string | query | False | Filter by rfc2307.member.uid |
| rfc2307.member.nis_netgroup | string | query | False | Filter by rfc2307.member.nis_netgroup |
| rfc2307.cn.group | string | query | False | Filter by rfc2307.cn.group |
| rfc2307.cn.netgroup | string | query | False | Filter by rfc2307.cn.netgroup |
| rfc2307.nis.object | string | query | False | Filter by rfc2307.nis.object |
| rfc2307.nis.mapname | string | query | False | Filter by rfc2307.nis.mapname |
| rfc2307.nis.netgroup_triple | string | query | False | Filter by rfc2307.nis.netgroup_triple |
| rfc2307.nis.netgroup | string | query | False | Filter by rfc2307.nis.netgroup |
| rfc2307.nis.mapentry | string | query | False | Filter by rfc2307.nis.mapentry |
| owner.name | string | query | False | Filter by owner.name |
| owner.uuid | string | query | False | Filter by owner.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| comment | string | query | False | Filter by comment |
| name_mapping.account.unix | string | query | False | Filter by name_mapping.account.unix |
| name_mapping.account.windows | string | query | False | Filter by name_mapping.account.windows |
| name_mapping.windows_to_unix.attribute | string | query | False | Filter by name_mapping.windows_to_unix.attribute |
| name_mapping.windows_to_unix.object_class | string | query | False | Filter by name_mapping.windows_to_unix.object_class |
| name_mapping.windows_to_unix.no_domain_prefix | boolean | query | False | Filter by name_mapping.windows_to_unix.no_domain_prefix |
| scope | string | query | False | Filter by scope |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of LDAP schema records. |
| records | array[ldap_schema] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "Schema based on Active Directory Services for UNIX
(read-only).",
      "global_schema": 1,
      "name": "AD-SFU-v1",
      "name_mapping": {
        "account": {
          "unix": "windowsAccount",
          "windows": "windowsAccount"
        },
        "windows_to_unix": {
          "attribute": "windowsAccount",
          "no_domain_prefix": "",
          "object_class": "User"
        }
      },
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "rfc2307": {
        "attribute": {
          "gecos": "name",
```

```
          "gid_number": "msSFU30GidNumber",
          "home_directory": "msSFU30HomeDirectory",
          "login_shell": "msSFU30LoginShell",
          "uid": "sAMAccountName",
          "uid_number": "msSFU30UidNumber",
          "user_password": "msSFU30Password"
        },
        "cn": {
          "group": "cn",
          "netgroup": "name"
        },
        "member": {
          "nis_netgroup": "msSFU30MemberNisNetgroup",
          "uid": "msSFU30MemberUid"
        },
        "nis": {
          "mapentry": "msSFU30NisMapEntry",
          "mapname": "msSFU30NisMapName",
          "netgroup": "msSFU30NisNetGroup",
          "netgroup_triple": "msSFU30MemberOfNisNetgroup",
          "object": "msSFU30NisObject"
        },
        "posix": {
          "account": "User",
          "group": "Group"
        }
      },
      "rfc2307bis": {
        "enabled": "",
        "group_of_unique_names": "groupOfUniqueNames",
        "maximum_groups": 256,
        "unique_member": "uniqueMember"
      },
      "scope": "string"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ldap_schema_account

| Name | Type | Description |
|---|---|---|
| unix | string | Attribute name used to retrieve UNIX account information. |
| windows | string | Attribute name used to retrieve Windows account information for a UNIX user account. |

windows_to_unix

| Name | Type | Description |
|---|---|---|
| attribute | string | Attribute name used to retrieve the UNIX account information for a Windows user account. |
| no_domain_prefix | boolean | Indicates whether or not the name for Windows to UNIX name mapping should have a domain prefix. |
| object_class | string | Name used to represent the windowsToUnix object class. |

ldap_schema_name_mapping

| Name | Type | Description |
| --- | --- | --- |
| account | ldap_schema_account | |
| windows_to_unix | windows_to_unix | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

attribute

| Name | Type | Description |
| --- | --- | --- |
| gecos | string | RFC 2307 gecos attribute. |
| gid_number | string | RFC 2307 gidNumber attribute. |
| home_directory | string | RFC 2307 homeDirectory attribute. |
| login_shell | string | RFC 2307 loginShell attribute. |
| uid | string | RFC 1274 userid attribute used by RFC 2307 as UID. |
| uid_number | string | RFC 2307 uidNumber attribute. |
| user_password | string | RFC 2256 userPassword attribute used by RFC 2307. |

cn

| Name | Type | Description |
| --- | --- | --- |
| group | string | RFC 2256 cn attribute used by RFC 2307 when working with groups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup | string | RFC 2256 cn attribute used by RFC 2307 when working with netgroups. |

member

| Name | Type | Description |
| --- | --- | --- |
| nis_netgroup | string | RFC 2307 memberNisNetgroup attribute. |
| uid | string | RFC 2307 memberUid attribute. |

nis

| Name | Type | Description |
| --- | --- | --- |
| mapentry | string | RFC 2307 nisMapEntry attribute. |
| mapname | string | RFC 2307 nisMapName attribute. |
| netgroup | string | RFC 2307 nisNetgroup object class. |
| netgroup_triple | string | RFC 2307 nisNetgroupTriple attribute. |
| object | string | RFC 2307 nisObject object class. |

posix

| Name | Type | Description |
| --- | --- | --- |
| account | string | RFC 2307 posixAccount object class. |
| group | string | RFC 2307 posixGroup object class. |

rfc2307

| Name | Type | Description |
| --- | --- | --- |
| attribute | attribute | |
| cn | cn | |

| Name | Type | Description |
| --- | --- | --- |
| member | member | |
| nis | nis | |
| posix | posix | |

rfc2307bis

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether RFC 2307bis is enabled for the client schema. |
| group_of_unique_names | string | RFC 2307bis groupOfUniqueNames object class. |
| maximum_groups | integer | Maximum number of groups supported when RFC 2307bis is enabled. |
| unique_member | string | RFC 2307bis uniqueMember attribute. |

template

The existing schema template you want to copy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the schema. |

ldap_schema

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | Comment to associate with the schema. |
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name | string | The name of the schema being created, modified or deleted. |
| name_mapping | ldap_schema_name_mapping | |

| Name | Type | Description |
|---|---|---|
| owner | owner | SVM, applies only to SVM-scoped objects. |
| rfc2307 | rfc2307 | |
| rfc2307bis | rfc2307bis | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an LDAP schema

POST `/name-services/ldap-schemas`

**Introduced In:** 9.11

Creates an LDAP schema.

#### Important notes

- To create a new schema, first create a copy of the default schemas provided by ONTAP and then modify the copy accordingly.
- If no value is specified for the owner.uuid or owner.name fields, the cserver UUID and name are used by default.

#### Related ONTAP commands

- `vserver services name-service ldap client schema copy`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name | string | The name of the schema being created, modified or deleted. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| template | template | The existing schema template you want to copy. |

### Example request

```
{
  "global_schema": 1,
  "name": "AD-SFU-v1",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "template": {
    "name": "AD-SFU-v1"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name | string | The name of the schema being created, modified or deleted. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| template | template | The existing schema template you want to copy. |

### Example response

```
{
  "global_schema": 1,
  "name": "AD-SFU-v1",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "template": {
    "name": "AD-SFU-v1"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name. |
| 4915221 | LDAP schema name in use in data SVM. |
| 4915222 | LDAP schema name in use in admin SVM. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ldap_schema_account

| Name | Type | Description |
|---|---|---|
| unix | string | Attribute name used to retrieve UNIX account information. |
| windows | string | Attribute name used to retrieve Windows account information for a UNIX user account. |

windows_to_unix

| Name | Type | Description |
|---|---|---|
| attribute | string | Attribute name used to retrieve the UNIX account information for a Windows user account. |
| no_domain_prefix | boolean | Indicates whether or not the name for Windows to UNIX name mapping should have a domain prefix. |
| object_class | string | Name used to represent the windowsToUnix object class. |

ldap_schema_name_mapping

| Name | Type | Description |
|---|---|---|
| account | ldap_schema_account | |
| windows_to_unix | windows_to_unix | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

attribute

| Name | Type | Description |
| --- | --- | --- |
| gecos | string | RFC 2307 gecos attribute. |
| gid_number | string | RFC 2307 gidNumber attribute. |
| home_directory | string | RFC 2307 homeDirectory attribute. |
| login_shell | string | RFC 2307 loginShell attribute. |
| uid | string | RFC 1274 userid attribute used by RFC 2307 as UID. |
| uid_number | string | RFC 2307 uidNumber attribute. |
| user_password | string | RFC 2256 userPassword attribute used by RFC 2307. |

cn

| Name | Type | Description |
| --- | --- | --- |
| group | string | RFC 2256 cn attribute used by RFC 2307 when working with groups. |
| netgroup | string | RFC 2256 cn attribute used by RFC 2307 when working with netgroups. |

member

| Name | Type | Description |
| --- | --- | --- |
| nis_netgroup | string | RFC 2307 memberNisNetgroup attribute. |
| uid | string | RFC 2307 memberUid attribute. |

nis

| Name | Type | Description |
| --- | --- | --- |
| mapentry | string | RFC 2307 nisMapEntry attribute. |
| mapname | string | RFC 2307 nisMapName attribute. |
| netgroup | string | RFC 2307 nisNetgroup object class. |
| netgroup_triple | string | RFC 2307 nisNetgroupTriple attribute. |
| object | string | RFC 2307 nisObject object class. |

posix

| Name | Type | Description |
| --- | --- | --- |
| account | string | RFC 2307 posixAccount object class. |
| group | string | RFC 2307 posixGroup object class. |

rfc2307

| Name | Type | Description |
| --- | --- | --- |
| attribute | attribute | |
| cn | cn | |
| member | member | |
| nis | nis | |
| posix | posix | |

rfc2307bis

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether RFC 2307bis is enabled for the client schema. |
| group_of_unique_names | string | RFC 2307bis groupOfUniqueNames object class. |
| maximum_groups | integer | Maximum number of groups supported when RFC 2307bis is enabled. |
| unique_member | string | RFC 2307bis uniqueMember attribute. |

template

The existing schema template you want to copy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the schema. |

ldap_schema

| Name | Type | Description |
| --- | --- | --- |
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name | string | The name of the schema being created, modified or deleted. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| template | template | The existing schema template you want to copy. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an existing schema

DELETE `/name-services/ldap-schemas/{owner.uuid}/{name}`

**Introduced In:** 9.11

Deletes an existing schema.

#### Related ONTAP commands

- `vserver services name-service ldap client schema delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| name | string | path | True | LDAP schema name. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4915204 | Failed to delete. The LDAP schema is being used by at least one LDAP client configuration. |
| 4915205 | The LDAP schema is a default schema and cannot be modified or deleted. |
| 4915217 | LDAP schema is owned by the admin SVM. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the LDAP schema details for an owner and schema

GET `/name-services/ldap-schemas/{owner.uuid}/{name}`

**Introduced In:** 9.11

Retrieves LDAP schema details for a given owner and schema.

#### Related ONTAP commands

- `vserver services name-service ldap client schema show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| name | string | path | True | LDAP schema name. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | Comment to associate with the schema. |
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name | string | The name of the schema being created, modified or deleted. |
| name_mapping | ldap_schema_name_mapping | |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| rfc2307 | rfc2307 | |
| rfc2307bis | rfc2307bis | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "Schema based on Active Directory Services for UNIX (read-only).",
  "global_schema": 1,
  "name": "AD-SFU-v1",
  "name_mapping": {
    "account": {
      "unix": "windowsAccount",
      "windows": "windowsAccount"
    },
    "windows_to_unix": {
      "attribute": "windowsAccount",
      "no_domain_prefix": "",
      "object_class": "User"
    }
  },
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "rfc2307": {
    "attribute": {
      "gecos": "name",
      "gid_number": "msSFU30GidNumber",
      "home_directory": "msSFU30HomeDirectory",
      "login_shell": "msSFU30LoginShell",
      "uid": "sAMAccountName",
      "uid_number": "msSFU30UidNumber",
      "user_password": "msSFU30Password"
    },
    "cn": {
      "group": "cn",
      "netgroup": "name"
    },
```

```
    "member": {
      "nis_netgroup": "msSFU30MemberNisNetgroup",
      "uid": "msSFU30MemberUid"
    },
    "nis": {
      "mapentry": "msSFU30NisMapEntry",
      "mapname": "msSFU30NisMapName",
      "netgroup": "msSFU30NisNetGroup",
      "netgroup_triple": "msSFU30MemberOfNisNetgroup",
      "object": "msSFU30NisObject"
    },
    "posix": {
      "account": "User",
      "group": "Group"
    }
  },
  "rfc2307bis": {
    "enabled": "",
    "group_of_unique_names": "groupOfUniqueNames",
    "maximum_groups": 256,
    "unique_member": "uniqueMember"
  },
  "scope": "string"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ldap_schema_account

| Name | Type | Description |
|---|---|---|
| unix | string | Attribute name used to retrieve UNIX account information. |
| windows | string | Attribute name used to retrieve Windows account information for a UNIX user account. |

windows_to_unix

| Name | Type | Description |
|---|---|---|
| attribute | string | Attribute name used to retrieve the UNIX account information for a Windows user account. |
| no_domain_prefix | boolean | Indicates whether or not the name for Windows to UNIX name mapping should have a domain prefix. |
| object_class | string | Name used to represent the windowsToUnix object class. |

ldap_schema_name_mapping

| Name | Type | Description |
|---|---|---|
| account | ldap_schema_account | |
| windows_to_unix | windows_to_unix | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

attribute

| Name | Type | Description |
| --- | --- | --- |
| gecos | string | RFC 2307 gecos attribute. |
| gid_number | string | RFC 2307 gidNumber attribute. |
| home_directory | string | RFC 2307 homeDirectory attribute. |
| login_shell | string | RFC 2307 loginShell attribute. |
| uid | string | RFC 1274 userid attribute used by RFC 2307 as UID. |
| uid_number | string | RFC 2307 uidNumber attribute. |
| user_password | string | RFC 2256 userPassword attribute used by RFC 2307. |

cn

| Name | Type | Description |
| --- | --- | --- |
| group | string | RFC 2256 cn attribute used by RFC 2307 when working with groups. |
| netgroup | string | RFC 2256 cn attribute used by RFC 2307 when working with netgroups. |

member

| Name | Type | Description |
| --- | --- | --- |
| nis_netgroup | string | RFC 2307 memberNisNetgroup attribute. |
| uid | string | RFC 2307 memberUid attribute. |

nis

| Name | Type | Description |
| --- | --- | --- |
| mapentry | string | RFC 2307 nisMapEntry attribute. |
| mapname | string | RFC 2307 nisMapName attribute. |
| netgroup | string | RFC 2307 nisNetgroup object class. |
| netgroup_triple | string | RFC 2307 nisNetgroupTriple attribute. |
| object | string | RFC 2307 nisObject object class. |

posix

| Name | Type | Description |
| --- | --- | --- |
| account | string | RFC 2307 posixAccount object class. |
| group | string | RFC 2307 posixGroup object class. |

rfc2307

| Name | Type | Description |
| --- | --- | --- |
| attribute | attribute | |
| cn | cn | |
| member | member | |
| nis | nis | |
| posix | posix | |

rfc2307bis

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether RFC 2307bis is enabled for the client schema. |
| group_of_unique_names | string | RFC 2307bis groupOfUniqueNames object class. |
| maximum_groups | integer | Maximum number of groups supported when RFC 2307bis is enabled. |
| unique_member | string | RFC 2307bis uniqueMember attribute. |

template

The existing schema template you want to copy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the schema. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the LDAP schema details for an owner and schema

PATCH `/name-services/ldap-schemas/{owner.uuid}/{name}`

**Introduced In:** 9.11

Updates LDAP schema details for a given owner and schema.

#### Important notes

- The default LDAP schemas provided by ONTAP cannot be modified.
- LDAP schemas can only be modified by the owner of the schema.

#### Related ONTAP commands

- `vserver services name-service ldap client schema modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| name | string | path | True | LDAP schema name. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | Comment to associate with the schema. |
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name_mapping | ldap_schema_name_mapping | |
| rfc2307 | rfc2307 | |
| rfc2307bis | rfc2307bis | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

**Example request**

```
{
  "comment": "Schema based on Active Directory Services for UNIX (read-only).",
  "global_schema": 1,
  "name_mapping": {
    "account": {
      "unix": "windowsAccount",
      "windows": "windowsAccount"
    },
    "windows_to_unix": {
      "attribute": "windowsAccount",
      "no_domain_prefix": "",
      "object_class": "User"
    }
  },
  "rfc2307": {
    "attribute": {
      "gecos": "name",
      "gid_number": "msSFU30GidNumber",
      "home_directory": "msSFU30HomeDirectory",
      "login_shell": "msSFU30LoginShell",
      "uid": "sAMAccountName",
      "uid_number": "msSFU30UidNumber",
      "user_password": "msSFU30Password"
    },
    "cn": {
      "group": "cn",
      "netgroup": "name"
    },
    "member": {
      "nis_netgroup": "msSFU30MemberNisNetgroup",
      "uid": "msSFU30MemberUid"
    },
    "nis": {
      "mapentry": "msSFU30NisMapEntry",
      "mapname": "msSFU30NisMapName",
      "netgroup": "msSFU30NisNetGroup",
      "netgroup_triple": "msSFU30MemberOfNisNetgroup",
      "object": "msSFU30NisObject"
    },
    "posix": {
      "account": "User",
      "group": "Group"
    }
```

```
    },
    "rfc2307bis": {
      "enabled": "",
      "group_of_unique_names": "groupOfUniqueNames",
      "maximum_groups": 256,
      "unique_member": "uniqueMember"
    },
    "scope": "string"
  }
```


#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4915205 | The LDAP schema is a default schema and cannot be modified or deleted. |
| 4915217 | LDAP schema is owned by the admin SVM. |
| 4915223 | LDAP schema does not belong to the admin SVM. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ldap_schema_account

| Name | Type | Description |
|---|---|---|
| unix | string | Attribute name used to retrieve UNIX account information. |
| windows | string | Attribute name used to retrieve Windows account information for a UNIX user account. |

windows_to_unix

| Name | Type | Description |
|---|---|---|
| attribute | string | Attribute name used to retrieve the UNIX account information for a Windows user account. |
| no_domain_prefix | boolean | Indicates whether or not the name for Windows to UNIX name mapping should have a domain prefix. |
| object_class | string | Name used to represent the windowsToUnix object class. |

ldap_schema_name_mapping

| Name | Type | Description |
|---|---|---|
| account | ldap_schema_account | |
| windows_to_unix | windows_to_unix | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

attribute

| Name | Type | Description |
| --- | --- | --- |
| gecos | string | RFC 2307 gecos attribute. |
| gid_number | string | RFC 2307 gidNumber attribute. |
| home_directory | string | RFC 2307 homeDirectory attribute. |
| login_shell | string | RFC 2307 loginShell attribute. |
| uid | string | RFC 1274 userid attribute used by RFC 2307 as UID. |
| uid_number | string | RFC 2307 uidNumber attribute. |
| user_password | string | RFC 2256 userPassword attribute used by RFC 2307. |

cn

| Name | Type | Description |
| --- | --- | --- |
| group | string | RFC 2256 cn attribute used by RFC 2307 when working with groups. |
| netgroup | string | RFC 2256 cn attribute used by RFC 2307 when working with netgroups. |

member

| Name | Type | Description |
| --- | --- | --- |
| nis_netgroup | string | RFC 2307 memberNisNetgroup attribute. |
| uid | string | RFC 2307 memberUid attribute. |

nis

| Name | Type | Description |
| --- | --- | --- |
| mapentry | string | RFC 2307 nisMapEntry attribute. |
| mapname | string | RFC 2307 nisMapName attribute. |
| netgroup | string | RFC 2307 nisNetgroup object class. |
| netgroup_triple | string | RFC 2307 nisNetgroupTriple attribute. |
| object | string | RFC 2307 nisObject object class. |

posix

| Name | Type | Description |
| --- | --- | --- |
| account | string | RFC 2307 posixAccount object class. |
| group | string | RFC 2307 posixGroup object class. |

rfc2307

| Name | Type | Description |
| --- | --- | --- |
| attribute | attribute | |
| cn | cn | |
| member | member | |
| nis | nis | |
| posix | posix | |

rfc2307bis

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Indicates whether RFC 2307bis is enabled for the client schema. |
| group_of_unique_names | string | RFC 2307bis groupOfUniqueNames object class. |
| maximum_groups | integer | Maximum number of groups supported when RFC 2307bis is enabled. |
| unique_member | string | RFC 2307bis uniqueMember attribute. |

template

The existing schema template you want to copy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the schema. |

ldap_schema

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Comment to associate with the schema. |
| global_schema | boolean | A global schema that can be used by all the SVMs. |
| name_mapping | ldap_schema_name_mapping | |
| rfc2307 | rfc2307 | |
| rfc2307bis | rfc2307bis | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage local mapping for hostnames

### Manage local mapping for hostnames

#### Overview

Displays and manages local mapping for hostnames.

#### Retrieving all hosts table entries

The local-hosts GET endpoint to retrieve all hosts table entries.

#### Examples

**Retrieving all the fields of all hosts table entries.**

```
# The API:
GET /api/name-services/local-hosts

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/local-
hosts?return_records=true&return_timeout=15&fields=*" -H "accept:
application/hal+json"

# The response:
{
 "records": [
{
```

```
  "owner": {
  "uuid": "6cdf045c-87ae-11eb-a56a-0050568e0287",
  "name": "svm1",
  "_links": {
     "self": {
     "href": "/api/svm/svms/6cdf045c-87ae-11eb-a56a-0050568e0287"
     }
   }
  },
  "address": "1.1.1.1",
  "hostname": "host.sales.foo.com",
  "aliases": [
          "host1.sales.foo.com",
          "host2.sakes.foo.com"
  ],
  "_links": {
   "self": {
   "href": "/api/name-services/local-hosts/6cdf045c-87ae-11eb-a56a-
0050568e0287/1.1.1.1"
  }
  }
  },
{
  "owner": {
  "uuid": "9d080552-7bff-11eb-a56a-0050568e0287",
  "name": "svm2",
  "_links": {
   "self": {
   "href": "/api/svm/svms/9d080552-7bff-11eb-a56a-0050568e0287"
   }
  }
  },
  "address": "2.2.2.2",
  "hostname": "samplehost2",
  "_links": {
  "self": {
  "href": "/api/name-services/local-hosts/9d080552-7bff-11eb-a56a-
0050568e0287/2.2.2.2"
  }
  }
  }
  ],
  "num_records": 2,
  "_links": {
  "self": {
  "href": "/api/name-services/local-
```

```
hosts?return_records=true&return_timeout=15&fields=*"
  }
  }
}
```

**Retrieving the hosts table entry of a given svm and address(ipv4/ipv6).**

```
# The API:
GET /api/name-services/local-hosts/{owner.uuid}/{address}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/local-hosts/9d080552-
7bff-11eb-a56a-0050568e0287/3.3.3.3" -H "accept: application/hal+json"

# The response:
{
 "owner": {
 "uuid": "9d080552-7bff-11eb-a56a-0050568e0287",
 "name": "svm1",
 "_links": {
"self": {
  "href": "/api/svm/svms/9d080552-7bff-11eb-a56a-0050568e0287"
 }
}
 },
 "address": "3.3.3.3",
 "hostname": "samplehost3",
 "aliases": [
 "host1.sales.foo.com",
 "host2.sakes.foo.com"
],
 "_links": {
 "self": {
 "href": "/api/name-services/local-hosts/9d080552-7bff-11eb-a56a-
0050568e0287/3.3.3.3"
 }
 }
}
```

#### Creating a hosts table entry

The local-hosts POST endpoint creates a new hosts table entry.

#### Examples

**Creating a hosts table entry with all fields.**

```
# The API:
POST /api/name-services/local-hosts

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/local-
hosts?return_records=false" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"address\": \"3.3.3.3\", \"aliases\": [
\"host1.sales.foo.com\", \"host2.sakes.foo.com\" ],
\"hostname\":\"samplehost3\", \"owner\": { \"name\": \"svm2\", \"uuid\":
\"9d080552-7bff-11eb-a56a-0050568e0287\" }}"
```

**Creating a hosts table entry with only required fields.**

```
# The API:
POST /api/name-services/local-hosts

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/local-
hosts?return_records=false" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"address\": \"123.123.123.12\",
\"hostname\": \"host.sales.foo.com\", \"owner\": { \"name\": \"svm2\",
\"uuid\": \"9d080552-7bff-11eb-a56a-0050568e0287\" }}"
```

#### Updating a hosts table entry

The local-hosts PATCH endpoint updates an existing hosts table entry.

#### Example

**Updating aliases and hostname of a given svm and address(ipv4/ipv6).**

```
# The API:
PATCH /api/name-services/local-hosts/{owner.uuid}/{address}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/local-hosts/9d080552-
7bff-11eb-a56a-0050568e0287/123.123.123.123" -H "accept: application/json"
-H "Content-Type: application/json" -d "{ \"aliases\": [
\"host1.sales.foo.com\", \"host2.sakes.foo.com\" ], \"hostname\":
\"host.sales.foo.com\"}"
```

---

#### Deleting a hosts table entry

---

The local-hosts DELETE endpoint deletes an existing hosts table entry.

##### Example

**Deleting the hosts table entry of a given svm and address(ipv4/ipv6).**

```
# The API:
DELETE /api/name-services/local-hosts/{owner.uuid}/{address}

#The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/local-hosts/9d080552-
7bff-11eb-a56a-0050568e0287/1.1.1.1" -H "accept: application/json"
```

---

### Retrieve IP to hostname mappings for all SVMs in a cluster

GET `/name-services/local-hosts`

**Introduced In:** 9.10

Retrieves all IP to hostname mappings for all SVMs of the cluster.

##### Related ONTAP commands

- `vserver services name-service dns hosts show`

##### Learn more

- DOC /name-services/local-hosts

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scope | string | query | False | Filter by scope |
| hostname | string | query | False | Filter by hostname<br>• maxLength: 255<br>• minLength: 1 |
| owner.name | string | query | False | Filter by owner.name |
| owner.uuid | string | query | False | Filter by owner.uuid |
| aliases | string | query | False | Filter by aliases |
| address | string | query | False | Filter by address |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of hosts to IP records. |
| records | array[local_host] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "address": "123.123.123.123",
      "aliases": [
        "host1.sales.foo.com",
        "host2.sakes.foo.com"
      ],
      "hostname": "host.sales.foo.com",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "string"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_host

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| address | string | IPv4/IPv6 address in dotted form. |
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| owner | owner | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
|---|---|---|
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a new IP to hostname mapping

POST `/name-services/local-hosts`

**Introduced In:** 9.10

Creates a new cluster-scoped or SVM-scoped IP to hostname mapping. For SVM-scoped mappings, include either the SVM name as owner.name or the SVM UUID as owner.uuid in the request body, along with other necessary parameters. For cluster-scoped mappings, specifying owner.uuid or owner.name is not required.

#### Required properties

- `address` - IPv4/IPv6 address in dotted form.
- `hostname` - Canonical hostname.

#### Optional properties

- `owner.uuid` or `owner.name` - Specify the name or UUID of an existing SVM to create an SVM-scoped IP-to-host mapping.
- `aliases` - The list of aliases.

#### Related ONTAP commands

- `vserver services name-service dns hosts create`

#### Learn more

- DOC /name-services/local-hosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| address | string | IPv4/IPv6 address in dotted form. |
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

### Example request

```
{
  "address": "123.123.123.123",
  "aliases": [
    "host1.sales.foo.com",
    "host2.sakes.foo.com"
  ],
  "hostname": "host.sales.foo.com",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1377682 | IPv6 is not enabled in the cluster. |
| 1966253 | IPv6 is not enabled in the cluster. Enable IPv6 using command "network options ipv6 modify -enabled true" and try again. |
| 2621706 | The specified owner UUID is incorrect for the specified owner name. |
| 8912896 | Only admin or data Vservers allowed. |

| Error Code | Description |
| --- | --- |
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |
| 23724155 | The specified IPv4 address is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast. |
| 23724156 | The specified IPv6 address is not supported because it is one of the following: ::, link-local, multicast, v4-compatible, v4-mapped, or loopback. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_host

| Name | Type | Description |
|---|---|---|
| address | string | IPv4/IPv6 address in dotted form. |
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an existing host object

DELETE `/name-services/local-hosts/{owner.uuid}/{address}`

**Introduced In:** 9.10

Deletes an existing host object.

#### Related ONTAP commands

- `vserver services name-service dns hosts delete`

#### Learn more

- DOC /name-services/local-hosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| address | string | path | True | The IP address. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a hostname mapping for an SVM

GET `/name-services/local-hosts/{owner.uuid}/{address}`

**Introduced In:** 9.10

For a specified SVM and IP address, returns the corresponding IP to hostname mapping.

#### Related ONTAP commands

- `vserver services name-service dns hosts show`

#### Learn more

- DOC /name-services/local-hosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| address | string | path | True | The IP address. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| address | string | IPv4/IPv6 address in dotted form. |
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| owner | owner | SVM, applies only to SVM-scoped objects. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "address": "123.123.123.123",
  "aliases": [
    "host1.sales.foo.com",
    "host2.sakes.foo.com"
  ],
  "hostname": "host.sales.foo.com",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string"
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a hostname mapping for an SVM

PATCH `/name-services/local-hosts/{owner.uuid}/{address}`

**Introduced In:** 9.10

For a specified SVM and IP address, modifies the corresponding IP to hostname mapping.

#### Related ONTAP commands

- `vserver services name-service dns hosts modify`

#### Learn more

- DOC /name-services/local-hosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | UUID of the owner to which this object belongs. |
| address | string | path | True | The IP address. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

#### Example request

```
{
  "aliases": [
    "host1.sales.foo.com",
    "host2.sakes.foo.com"
  ],
  "hostname": "host.sales.foo.com",
  "scope": "string"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724055 | Internal error. Configuration for Vserver failed. Verify that the cluster is healthy, then try the command again. For further assistance, contact technical support. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

owner

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

local_host

| Name | Type | Description |
|---|---|---|
| aliases | array[string] | The list of aliases. |
| hostname | string | Canonical hostname. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage name mappings for SVMs

### Manage name mappings for SVMs

#### Overview

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, and UNIX identities to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client or a CIFS client.

The system keeps a set of conversion rules for each Storage Virtual Machine (SVM). Each rule consists of two pieces: a pattern and a replacement. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX sed program.

Name mappings are applied in the order in which they occur in the priority list; for example, a name mapping that occurs at position 2 in the priority list is applied before a name mapping that occurs at position 3. Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX, and UNIX-to-Windows) has its own priority list. You are prevented from creating two name mappings with the same pattern.

#### Examples

### Creating a name-mapping with client_match as the ip-address

Use the following API to create a name-mapping. Note the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /api//name-services/name-mappings

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/name-
mappings?return_records=true" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"client_match\": \"10.254.101.111/28\",
\"direction\": \"win_unix\", \"index\": 1, \"pattern\":
\"ENGCIFS_AD_USER\", \"replacement\": \"unix_user1\", \"svm\": { \"name\":
\"vs1\", \"uuid\": \"f71d3640-0226-11e9-8526-000c290a8c4b\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
      "name": "vs1"
    },
    "direction": "win_unix",
    "index": 1,
    "pattern": "ENGCIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "10.254.101.111/28"
  }
]
}
```

#### Creating a name-mapping with client_match as the hostname

Use the following API to create a name-mapping. Note the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /api//name-services/name-mappings

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/name-
mappings?return_records=true" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"client_match\": \"google.com\",
\"direction\": \"win_unix\", \"index\": 2, \"pattern\":
\"ENGCIFS_AD_USER\", \"replacement\": \"unix_user1\", \"svm\": { \"name\":
\"vs1\", \"uuid\": \"f71d3640-0226-11e9-8526-000c290a8c4b\" }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
      "name": "vs1"
    },
    "direction": "win_unix",
    "index": 2,
    "pattern": "ENGCIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "google.com"
  }
]
}
```

**Retrieving all name-mapping configurations for all SVMs in the cluster**

```
# The API:
GET /api/name-services/name-mappings

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/name-
mappings?fields=*&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
      "name": "vs1"
    },
    "direction": "win_unix",
    "index": 1,
    "pattern": "ENGCIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "10.254.101.111/28"
  },
  {
    "svm": {
      "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
      "name": "vs1"
    },
    "direction": "win_unix",
    "index": 2,
    "pattern": "ENGCIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "google.com"
  }
],
"num_records": 2
}
```

**Retrieving a name-mapping configuration for a specific SVM, and for the specified direction and index**

```
# The API:
GET /api/name-services/name-mappings/{svm.uuid}/{direction}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-
0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
  "name": "vs1"
},
"direction": "win_unix",
"index": 1,
"pattern": "ENGCIFS_AD_USER",
"replacement": "unix_user1",
"client_match": "10.254.101.111/28"
}
```

#### Updating a specific name-mapping configuration

```
# The API:
PATCH /api//name-services/name-mappings/{svm.uuid}/{direction}/{index}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-
0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"client_match\":
\"10.254.101.222/28\", \"pattern\": \"ENGCIFS_LOCAL_USER\",
\"replacement\": \"pcuser\"}"

# swapping a specified namemapping entry by index
curl -X PATCH "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-
0226-11e9-8526-000c290a8c4b/win-unix/3?new_index=1" -H "accept:
application/json" -H "Content-Type: application/json" -d "{  \"pattern\":
\"ENGCIFS_AD_USER\", \"replacement\": \"unix_user1\"}"
```

### Removing a specific name-mapping configuration

```
# The API:
DELETE /api/name-services/name-mappings/{svm.uuid}/{direction}/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/name-
mappings/f71d3640-0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept:
application/json"
```

## Retrieve the hostname mapping for all SVMs

GET `/name-services/name-mappings`

**Introduced In:** 9.6

Retrieves the name mapping configuration for all SVMs.

#### Related ONTAP commands

- `vserver name-mapping show`

#### Learn more

- DOC /name-services/name-mappings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| replacement | string | query | False | Filter by replacement<br>• maxLength: 256<br>• minLength: 1 |
| index | integer | query | False | Filter by index<br>• Max value: 2147483647<br>• Min value: 1 |
| client_match | string | query | False | Filter by client_match |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| pattern | string | query | False | Filter by pattern<br>• maxLength: 256<br>• minLength: 1 |
| direction | string | query | False | Filter by direction |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[name_mapping] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_match": "10.254.101.111/28",
      "direction": "win_unix",
      "index": 1,
      "pattern": "ENGCIFS_AD_USER",
      "replacement": "unix_user1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

name_mapping

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, UNIX identities to CIFS identities, S3 to UNIX identities and S3 to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client, CIFS client or an S3 client.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |
| direction | string | Direction in which the name mapping is applied. The possible values are:<br>• krb_unix - Kerberos principal name to UNIX user name<br>• win_unix - Windows user name to UNIX user name<br>• unix_win - UNIX user name to Windows user name mapping<br>• s3_unix - S3 user name to UNIX user name mapping<br>• s3_win - S3 user name to Windows user name mapping |
| index | integer | Position in the list of name mappings. |

| Name | Type | Description |
| --- | --- | --- |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create hostname mappings for an SVM**

POST `/name-services/name-mappings`

**Introduced In:** 9.6

Creates name mappings for an SVM.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the name mapping.
- `index` - Name mapping's position in the priority list.
- `direction` - Direction of the name mapping.
- `pattern` - Pattern to match to. Maximum length is 256 characters.
- `replacement` - Replacement pattern to match to. Maximum length is 256 characters.

#### Recommended optional properties

- `client_match` - Hostname or IP address added to match the pattern to the client's workstation IP address.

#### Related ONTAP commands

- `vserver name-mapping create`
- `vserver name-mapping insert`

#### Learn more

- DOC /name-services/name-mappings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| insert_name_map | boolean | query | False | Create a name mapping at a specified position<br>• Introduced in: 9.17<br>• Default value: |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |
| direction | string | Direction in which the name mapping is applied. The possible values are:<br>• krb_unix - Kerberos principal name to UNIX user name<br>• win_unix - Windows user name to UNIX user name<br>• unix_win - UNIX user name to Windows user name mapping<br>• s3_unix - S3 user name to UNIX user name mapping<br>• s3_win - S3 user name to Windows user name mapping |
| index | integer | Position in the list of name mappings. |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |

| Name | Type | Description |
| --- | --- | --- |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "client_match": "10.254.101.111/28",
  "direction": "win_unix",
  "index": 1,
  "pattern": "ENGCIFS_AD_USER",
  "replacement": "unix_user1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[name_mapping] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "client_match": "10.254.101.111/28",
      "direction": "win_unix",
      "index": 1,
      "pattern": "ENGCIFS_AD_USER",
      "replacement": "unix_user1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65798185 | Failed to resolve the specified hostname |
| 65798173 | The name mapping pattern already exists with a more generic IP qualifier |
| 65798149 | Invalid index for the name mapping entry |
| 2621706 | The specified svm.uuid and svm.name refer to different SVMs |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

name_mapping

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, UNIX identities to CIFS identities, S3 to UNIX identities and S3 to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client, CIFS client or an S3 client.

| Name | Type | Description |
|---|---|---|
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |

| Name | Type | Description |
| --- | --- | --- |
| direction | string | Direction in which the name mapping is applied. The possible values are:<br>• krb_unix - Kerberos principal name to UNIX user name<br>• win_unix - Windows user name to UNIX user name<br>• unix_win - UNIX user name to Windows user name mapping<br>• s3_unix - S3 user name to UNIX user name mapping<br>• s3_win - S3 user name to Windows user name mapping |
| index | integer | Position in the list of name mappings. |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the name mapping configuration

DELETE `/name-services/name-mappings/{svm.uuid}/{direction}/{index}`

**Introduced In:** 9.6

Deletes the name mapping configuration.

#### Related ONTAP commands

- `vserver name-mapping delete`

#### Learn more

- DOC /name-services/name-mappings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| direction | string | path | True | Direction |
| index | integer | path | True | Position of the entry in the list |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the name mapping configuration for an SVM

GET `/name-services/name-mappings/{svm.uuid}/{direction}/{index}`

**Introduced In:** 9.6

Retrieves the name mapping configuration of an SVM.

#### Related ONTAP commands

- `vserver name-mapping show`

#### Learn more

- DOC /name-services/name-mappings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| direction | string | path | True | Direction |
| index | integer | path | True | Position of the entry in the list |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |

| Name | Type | Description |
| --- | --- | --- |
| direction | string | Direction in which the name mapping is applied. The possible values are:<br>• krb_unix - Kerberos principal name to UNIX user name<br>• win_unix - Windows user name to UNIX user name<br>• unix_win - UNIX user name to Windows user name mapping<br>• s3_unix - S3 user name to UNIX user name mapping<br>• s3_win - S3 user name to Windows user name mapping |
| index | integer | Position in the list of name mappings. |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "client_match": "10.254.101.111/28",
  "direction": "win_unix",
  "index": 1,
  "pattern": "ENGCIFS_AD_USER",
  "replacement": "unix_user1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the name mapping configuration for an SVM

PATCH `/name-services/name-mappings/{svm.uuid}/{direction}/{index}`

**Introduced In:** 9.6

Updates the name mapping configuration of an SVM. The positions can be swapped by providing the `new_index` property.
Swapping is not allowed for entries that have `client_match` property configured.

#### Related ONTAP commands

- `vserver name-mapping modify`
- `vserver name-mapping swap`

#### Learn more

- DOC /name-services/name-mappings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| direction | string | path | True | Direction |
| index | integer | path | True | Position of the entry in the list |
| new_index | integer | query | False | New position of the Index after a swap is completed.<br>• Introduced in: 9.7 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |

**Example request**

```
{
  "client_match": "10.254.101.111/28",
  "pattern": "ENGCIFS_AD_USER",
  "replacement": "unix_user1"
}
```

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65798185 | Failed to resolve the specified hostname |
| 65798179 | Cannot swap entries because one or both entries have host name or address configured. |
| | Delete and re-create the new entry at the specified position. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

name_mapping

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, UNIX identities to CIFS identities, S3 to UNIX identities and S3 to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client, CIFS client or an S3 client.

| Name | Type | Description |
|---|---|---|
| client_match | string | Client workstation IP Address which is matched when searching for the pattern.<br>You can specify the value in any of the following formats:<br>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24<br>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64<br>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0<br>• As a hostname |

| Name | Type | Description |
| --- | --- | --- |
| pattern | string | Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows. |
| replacement | string | The name that is used as a replacement, if the pattern associated with this entry matches. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Display netgroup file details

### Display netgroup file details

#### Overview

Displays the netgroup file details or raw netgroup file of an SVM.
Note: The GET collection endpoint is not supported for netgroup files.

#### Examples

### Retrieving the netgroup file status of a given SVM

```
# The API:
GET /api/name-services/netgroup-files/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/netgroup-files/53942195-
1709-11ec-b0d4-0050568efd14" -H "accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "53942195-1709-11ec-b0d4-0050568efd14",
  "name": "svm1"
},
"timestamp": "2021-10-04T13:05:06+05:30",
"hash_value": "729b6e43cee04bcee18efa3aa689881d",
"hash_value_by_host": "eb109a44056a47bdeb4b407ec821a14b",
"file_size": 5118686
}
```

### Retrieving the raw netgroup file of a given SVM

```
# The API:
GET /api/name-services/netgroup-files/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/netgroup-files/53942195-
1709-11ec-b0d4-0050568efd14" -H "accept: multipart/form-data"
```

### Deleting a netgroup file of a given SVM

```
# The API:
DELETE /api/name-services/netgroup-files/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/netgroup-
files/53942195-1709-11ec-b0d4-0050568efd14" -H "accept: application/json"
```

---

### Delete a netgroup file for an SVM

DELETE `/name-services/netgroup-files/{svm.uuid}`

**Introduced In:** 9.11

Deletes the netgroup file for a given SVM.

#### Related ONTAP commands

- `vserver services name-service netgroup file delete`

#### Learn more

- DOC /name-services/netgroup-files/{svm.uuid}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the netgroup file details

GET `/name-services/netgroup-files/{svm.uuid}`

**Introduced In:** 9.11

Retrieves the netgroup file details for a given SVM if the header is set as "accept: application/json".
If the header is set as "accept: multipart/form-data", the raw netgroup file of the given SVM is retrieved.

#### Important notes

- For a raw netgroup file, set the header as "accept: multipart/form-data" else "accept: application/json" for netgroup file details.
- Maximum size supported for raw netgroup file is 5MB, but netgroup file details of any valid SVM can be retrieved if present.

#### Related ONTAP commands

- `vserver services name-service netgroup status`
- `vserver services name-service netgroup file show`

#### Learn more

- DOC /name-services/netgroup-files/{svm.uuid}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| file_size | integer | File size of the netgroup definitions file in bytes. |
| hash_value | string | Hash value of the netgroup definitions. |
| hash_value_by_host | string | Hash value of the netgroup-by-host database. |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| timestamp | string | Load time for netgroup definitions. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "file_size": 5118686,
  "hash_value": "e53ec87782356bd6786f146ce0a48449",
  "hash_value_by_host": "e012b7f62e4810936725ed1239018314",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timestamp": "string"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage NIS server configurations

## Manage NIS server configurations

### Overview

NIS servers are used to authenticate user and client computers. NIS domain name and NIS server information is required to configure NIS.
It is important to note that this API is used to retrieve and manage NIS server configurations for data SVMs only. NIS configuration for the cluster is managed via /api/security/authentication/cluster/nis .

### Retrieving NIS Information

The NIS GET endpoint retrieves all of the NIS configurations for data SVMs.

### Examples

## Retrieving all fields for all NIS configurations

```
# The API:
/api/name-services/nis

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/nis?fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domain": "domainA.example.com",
    "servers": [
      "10.10.10.10",
      "example.com"
    ],
    "bound-servers": [
      "10.10.10.10"
    ],
```

```
    "_links": {
      "self": {
        "href": "/api/name-services/nis/179d3c85-7053-11e8-b9b8-
005056b41bd1"
      }
    }
  },
  {
    "svm": {
      "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domain": "domainB.example.com",
    "servers": [
      "2.2.2.2",
      "3.3.3.3",
      "4.4.4.4"
    ],
    "bound-servers": [],
    "_links": {
      "self": {
        "href": "/api/name-services/nis/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/name-services/nis?fields=*"
  }
}
}
```


### Retrieving all NIS configurations whose bound servers start with *10*

```
# The API:
/api/name-services/nis

# The call:
curl -X GET "https://<mgmt-ip/api/name-services/nis?bound_servers=10*" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "bound-servers": [
      "10.10.10.10"
    ],
    "_links": {
      "self": {
        "href": "/api/name-services/nis/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/name-services/nis?bound_servers=10*"
  }
}
}
```

### Retrieving the NIS configuration of a specific SVM

```
# The API:
/api/name-services/nis/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1"
},
"domain": "domainA.example.com",
"servers": [
  "10.10.10.10",
  "example.com"
],
"bound_servers": [
  "10.10.10.10"
]
}
```

#### Creating a NIS configuration

The NIS POST endpoint creates a NIS configuration for the specified SVM.

#### Example

The following example shows a POST operation:

```
# The API:
/api/name-services/nis

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/nis" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domain\":
\"domainA.example.com\", \"servers\": [ \"10.10.10.10\",\"example.com\"
]}"
```

#### Updating the NIS configuration

The NIS PATCH endpoint updates the NIS configuration for the specified NIS server.

#### Examples

### Updating the domain

```
# The API:
/api/name-services/nis/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"domain\": \"domainC.example.com\", \"servers\":
[ \"13.13.13.13\" ]}"
```

### Updating the server

```
# The API:
/api/name-services/nis/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"servers\": [ \"14.14.14.14\" ]}"
```

#### Deleting a NIS configuration

The NIS DELETE endpoint deletes the NIS configuration for the specified SVM.

#### Example

The following example shows a DELETE operation:

```
# The API:
/api/name-services/nis/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-
11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"
```

---

### Retrieve NIS domain configurations of all SVMs

GET `/name-services/nis`

**Introduced In:** 9.6

Retrieves NIS domain configurations of all the SVMs. The bound_servers field indicates the successfully bound NIS servers. Lookups and authentications fail if there are no bound servers.

#### Related ONTAP commands

- `vserver services name-service nis-domain show`
- `vserver services name-service nis-domain show-bound`
- `vserver services name-service nis-domain show-bound-debug`

#### Learn more

- DOC /name-services/nis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| servers | string | query | False | Filter by servers<br>• maxLength: 255<br>• minLength: 1<br>• Introduced in: 9.7 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.7 |
| svm.uuid | string | query | False | Filter by svm.uuid<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| binding_details.server | string | query | False | Filter by binding_details.server<br>• maxLength: 255<br>• minLength: 1<br>• Introduced in: 9.11 |
| binding_details.status.code | string | query | False | Filter by binding_details.status.code<br>• Introduced in: 9.11 |
| binding_details.status.message | string | query | False | Filter by binding_details.status.message<br>• Introduced in: 9.11 |
| domain | string | query | False | Filter by domain<br>• maxLength: 64<br>• minLength: 1<br>• Introduced in: 9.7 |
| bound_servers | string | query | False | Filter by bound_servers<br>• maxLength: 255<br>• minLength: 1<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of NIS domain records. |
| records | array[nis_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "binding_details": [
        {
          "server": "string",
          "status": {
            "code": "string",
            "message": "string"
          }
        }
      ],
      "bound_servers": [
        "string"
      ],
      "domain": "domainA.example.com",
      "servers": [
        "10.10.10.10",
        "example.com"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
```

```
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nis_service

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the NIS domain and server configuration for a data SVM

POST `/name-services/nis`

**Introduced In:** 9.6

Creates an NIS domain and server configuration for a data SVM.
NIS configuration for the cluster is managed via /api/security/authentication/cluster/nis .

#### Important notes

- Each SVM can have one NIS domain configuration.
- Multiple SVMs can be configured with the same NIS domain. Specify the NIS domain and NIS servers as input.Domain name and servers fields cannot be empty.
- Both FQDNs and IP addresses are supported for the servers field.
- IPv6 must be enabled if IPv6 family addresses are specified in the servers field.
- A maximum of ten NIS servers are supported.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the NIS configuration.
- `domain` - NIS domain to which the configuration belongs.
- `servers` - List of NIS server IP addresses.

#### Related ONTAP commands

- `vserver services name-service nis-domain create`

#### Learn more

- DOC /name-services/nis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "domainA.example.com",
  "servers": [
    "10.10.10.10",
    "example.com"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of NIS domain records. |
| records | array[nis_service] | |

#### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "binding_details": [
        {
          "server": "string",
          "status": {
            "code": "string",
            "message": "string"
          }
        }
      ],
      "bound_servers": [
        "string"
      ],
      "domain": "domainA.example.com",
      "servers": [
        "10.10.10.10",
        "example.com"
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1966253 | IPv6 is not enabled in the cluster |
| 2621488 | Invalid SVM context |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 3276964 | NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters |
| 3276933 | A maximum of 10 NIS servers can be configured per SVM |
| 13434916 | The SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 23724109 | DNS resolution failed for one or more specified servers |
| 23724111 | Empty NIS servers cannot be specified. |
| 23724112 | DNS resolution failed due to an internal error. Contact technical support if this issue persists |
| 23724132 | DNS resolution failed for all the specified servers |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nis_service

| Name | Type | Description |
| --- | --- | --- |
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete the NIS domain configuration for an SVM**

DELETE `/name-services/nis/{svm.uuid}`

**Introduced In:** 9.6

Deletes the NIS domain configuration of an SVM. NIS can be removed as a source from ns-switch if NIS is not used for lookups.

#### Related ONTAP commands

- `vserver services name-service nis-domain delete`

#### Learn more

- DOC /name-services/nis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the NIS domain and server configurations for an SVM

GET `/name-services/nis/{svm.uuid}`

**Introduced In:** 9.6

Retrieves NIS domain and server configurations of an SVM. Both NIS domain and servers are displayed by default. The bound_servers field indicates the successfully bound NIS servers.

#### Related ONTAP commands

- `vserver services name-service nis-domain show`
- `vserver services name-service nis-domain show-bound`
- `vserver services name-service nis-domain show-bound-debug`

#### Learn more

- DOC /name-services/nis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "domainA.example.com",
  "servers": [
    "10.10.10.10",
    "example.com"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the NIS domain and server configuration for an SVM

PATCH `/name-services/nis/{svm.uuid}`

**Introduced In:** 9.6

Updates NIS domain and server configuration of an SVM.

#### Important notes

- Both NIS domain and servers can be modified.
- Domains and servers cannot be empty.
- Both FQDNs and IP addresses are supported for the servers field.
- If the domain is modified, NIS servers must also be specified.
- IPv6 must be enabled if IPv6 family addresses are specified for the servers field.

#### Related ONTAP commands

- `vserver services name-service nis-domain modify`

#### Learn more

- DOC /name-services/nis

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used<br>by the NIS domain configuration. |

**Example request**

```
{
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "domainA.example.com",
  "servers": [
    "10.10.10.10",
    "example.com"
  ]
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966253 | IPv6 is not enabled in the cluster |
| 2621488 | Invalid SVM context |
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 3276964 | NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters |
| 3276933 | A maximum of 10 NIS servers can be configured per SVM |
| 23724109 | DNS resolution failed for one or more specified servers |
| 23724111 | Empty NIS servers cannot be specified. |
| 23724112 | DNS resolution failed due to an internal error. Contact technical support if this issue persists |
| 23724132 | DNS resolution failed for all the specified servers |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 LIFs |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

nis_service

| Name | Type | Description |
| --- | --- | --- |
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage UNIX group configurations

### Manage UNIX group configurations

#### Overview

You can use this API to display local UNIX group information and to control UNIX group configurations.

#### Retrieving UNIX group information

The UNIX group GET endpoint retrieves all of the local UNIX groups configurations for data SVMs.

#### Examples

### Retrieving all of the fields for all of the UNIX group configurations

The UNIX group GET endpoint retrieves all of the local UNIX groups configurations for data SVMs.

```
# The API:
/api/name-services/unix-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-groups?fields=*" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "name": "group1",
    "id": 11,
    "users": [
      {
        "name": "user1"
      },
      {
        "name": "user2"
      },
      {
        "name": "user3"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/name-services/unix-groups/b009a9e7-4081-b576-7575-
ada21efcaf16/group1"
      }
    }
  },
  {
```

```
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "name": "group2",
    "id": 12,
    "users": [
      {
        "name": "user1"
      },
      {
        "name": "user2"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/name-services/unix-groups/b009a9e7-4081-b576-7575-
ada21efcaf16/group2"
      }
    }
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcad17",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcad17"
        }
      }
    },
    "name": "group1",
    "id": 11,
    "users": [
      {
        "name": "user2"
      },
      {
        "name": "user3"
      }
    ],
```

```
    "_links": {
      "self": {
        "href": "/api/name-services/unix-groups/b009a9e7-4081-b576-7575-
ada21efcad17/group1"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/name-services/unix-groups?fields=*"
  }
}
}
```


**Retrieving all of the UNIX group configurations whose group name is 'group1'.**

```
# The API:
/api/name-services/unix-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-groups?name=group1"
-H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "name": "group1",
    "id": 11,
    "_links": {
      "self": {
        "href": "/api/name-services/unix-groups/b009a9e7-4081-b576-7575-
ada21efcaf16/group1"
      }
    }
  },
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/name-services/unix-groups?name=group1"
  }
}
}
```

#### Creating a UNIX group configuration

The UNIX group POST endpoint creates a UNIX group configuration for the specified SVM.

#### Example

```
# The API:
/api/name-services/unix-groups

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/unix-groups" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d '{ "svm":
{"uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1"}, "name": "group1", "id":
111}'
```

#### Updating a UNIX group configuration

The UNIX group PATCH endpoint updates the UNIX group ID of the specified UNIX group and the specified SVM.

#### Example

### Modify the group ID of group1 to 112

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/unix-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/group1" -H "accept: application/hal+json" -H
"Content-Type: application/json" -d '{ "id": 112}'
```

#### Deleting a UNIX group configuration

The UNIX group DELETE endpoint deletes the specified UNIX group of the specified SVM.

#### Example

### Delete the group 'group1'

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/unix-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/group1" -H "accept: application/hal+json"
```

### Retrieve UNIX groups for all SVMs

GET `/name-services/unix-groups`

**Introduced In:** 9.9

Retrieves the UNIX groups for all of the SVMs. UNIX users who are the members of the group are also displayed.

#### Related ONTAP commands

- `vserver services name-service unix-group show`

#### Learn more

- DOC /name-services/unix-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| users.name | string | query | False | Filter by users.name |
| id | integer | query | False | Filter by id |
| name | string | query | False | Filter by name |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| current_count | integer | Number of UNIX groups and members. |
| num_records | integer | Number of UNIX group records. |
| records | array[unix_group] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "group1",
      "skip_name_validation": true,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        {
          "name": "string"
        }
      ]
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

unix_group

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| id | integer | UNIX group ID of the specified user. |
| name | string | UNIX group name to be added to the local database. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a local UNIX group configuration for an SVM**

POST `/name-services/unix-groups`

**Introduced In:** 9.9

Creates the local UNIX group configuration for the specified SVM.

Group name and group ID are mandatory parameters.

**Important notes**

- The default limit for local UNIX groups and group members is 32768.

#### Learn more

- DOC /name-services/unix-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| id | integer | UNIX group ID of the specified user. |
| name | string | UNIX group name to be added to the local database. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX group name is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | |

### Example request

```
{
  "name": "group1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "string"
    }
  ]
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name. |
| 2621516 | This operation is only supported on a data SVM. |
| 3277025 | Maximum supported limit of UNIX group count reached. |
| 3277051 | Invalid characters in group name. Valid characters are 0-9, A-Z, a-z, ".", "_" and "-". Names cannot start with "-". |

| Error Code | Description |
| --- | --- |
| 23724067 | Group name too long. Maximum supported length is 64 characters. |
| 23724141 | Duplicate group ID. Group ID must be unique. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

unix_group

| Name | Type | Description |
|---|---|---|
| id | integer | UNIX group ID of the specified user. |
| name | string | UNIX group name to be added to the local database. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX group name is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a UNIX group configuration for an SVM

DELETE `/name-services/unix-groups/{svm.uuid}/{name}`

**Introduced In:** 9.9

Deletes a UNIX group configuration for the specified SVM.

#### Related ONTAP commands

- `vserver services name-service unix-group delete`

#### Learn more

- DOC /name-services/unix-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve UNIX group information for a group and SVM

GET `/name-services/unix-groups/{svm.uuid}/{name}`

**Introduced In:** 9.9

Retrieves UNIX group information for the specified group and SVM. UNIX users who are part of this group are also retrieved.

#### Related ONTAP commands

- `vserver services name-service unix-group show`

#### Learn more

- DOC /name-services/unix-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| id | integer | UNIX group ID of the specified user. |
| name | string | UNIX group name to be added to the local database. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "group1",
  "skip_name_validation": true,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "string"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update UNIX group information for a group in an SVM

PATCH `/name-services/unix-groups/{svm.uuid}/{name}`

**Introduced In:** 9.9

Updates the UNIX group information of the specified group in the specified SVM.

#### Learn more

- DOC /name-services/unix-groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX group name. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| id | integer | UNIX group ID of the specified user. |
| users | array[users] | |

### Example request

```
{
  "skip_name_validation": true,
  "users": [
    {}
  ]
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 23724141 | Duplicate group ID. Group ID must be unique. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

unix_group

| Name | Type | Description |
|---|---|---|
| id | integer | UNIX group ID of the specified user. |
| users | array[users] | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage UNIX group users

### Manage UNIX group users

#### Overview

You can use this API to retrieve, add or delete UNIX users for an SVM's UNIX group.

#### Retrieving UNIX users from an SVM's UNIX group

The UNIX group users GET endpoint retrieves UNIX users of the specified UNIX group and the SVM.

#### Examples

### Retrieving all users from the group 'pcuser'

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-groups/f06686a2-
c901-11eb-94b4-0050568e9f2c/pcuser/users" -H "accept:
application/hal+json"

# The response:
{
 "records": [
 {
  "svm": {
    "uuid": "f06686a2-c901-11eb-94b4-0050568e9f2c",
    "name": "svm1"
  },
  "unix_group": {
    "name": "pcuser"
  },
  "name": "user1",
  "_links": {
    "self": {
      "href": "/api/name-services/unix-groups/f06686a2-c901-11eb-94b4-
0050568e9f2c/pcuser/users/user1"
    }
  }
 }
],
 "num_records": 1,
 "_links": {
 "self": {
   "href": "/api/name-services/unix-groups/f06686a2-c901-11eb-94b4-
0050568e9f2c/pcuser/users/?fields=*"
 }
}
}
```

**Retrieving user 'user1' from the group 'pcuser'**

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-groups/f06686a2-
c901-11eb-94b4-0050568e9f2c/pcuser/users/user1" -H "accept:
application/hal+json"

# The response:
{
 "svm": {
 "uuid": "f06686a2-c901-11eb-94b4-0050568e9f2c",
 "name": "svm1"
 },
 "unix_group": {
 "name": "pcuser"
 },
 "name": "user1",
 "_links": {
 "self": {
   "href": "/api/name-services/unix-groups/f06686a2-c901-11eb-94b4-
0050568e9f2c/pcuser/users/user1"
 }
}
}
```

#### Adding users to a UNIX group

The UNIX group users POST endpoint adds UNIX users to the specified UNIX group and the SVM. Multiple users can be added in a single call using the "records" parameter.

#### Examples

### Adding a single user to the group 'group1'

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/unix-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/group1/users" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d '{ "name":
"user4" }'
```

### Adding multiple users to the group 'group1' in a single REST call

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/unix-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/group1/users" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d '{ "records":
[{"name":"user1"}, {"name":"user2"}, {"name":"user3"}]}'
```

#### Deleting a user from a group of a specific SVM

#### Example

### Delete the user 'user1' from group 'group1' in SVM 'vs1'

```
# The API:
/api/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/unix-groups/179d3c85-
7053-11e8-b9b8-005056b41bd1/group1/users/user1" -H "accept:
application/hal+json"
```

## Retrieve users for a UNIX group and SVM

GET `/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users`

**Introduced In:** 9.10

Retrieves users for the specified UNIX group and SVM.

#### Related ONTAP commands

- `vserver services name-service unix-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_group.name | string | path | True | UNIX group name. |
| name | string | query | False | Filter by name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of UNIX user records in a given SVM and group. |
| records | array[unix_group_users] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "skip_name_validation": true,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_group": {
        "name": "string"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

records

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_group

| Name | Type | Description |
|---|---|---|
| name | string | UNIX group name. |

unix_group_users

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_group | unix_group | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add users to a UNIX group and SVM

POST `/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users`

**Introduced In:** 9.9

Adds users to the specified UNIX group and SVM.

#### Important notes

- Multiple users can be added in a single call using the "records" parameter.
- "records" parameter must not be specified when "name" parameter is specified.
- Specified users are appended to the existing list of users.

- Duplicate users are ignored.

#### Related ONTAP commands

- `vserver services name-service unix-group adduser`
- `vserver services name-service unix-group addusers`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_group.name | string | path | True | UNIX group name. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |
| records | array[records] | An array of UNIX users specified to add multiple users to a UNIX group in a single API call.<br>Not allowed when the `name` property is used. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX user names is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_group | unix_group | |

### Example request

```
{
  "name": "string",
  "records": [
    {
      "name": "string"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_group": {
    "name": "string"
  }
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262179 | Invalid arguments in records. |
| 3276897 | The specified UNIX group does not exist in the SVM. |
| 3277025 | Maximum supported limit of UNIX group membership count reached. |

| Error Code | Description |
| --- | --- |
| 3277051 | Invalid characters in user name. Valid characters are 0-9, A-Z, a-z, ".", "_" and "-". Names cannot start with "-". |
| 23724066 | User name too long. Maximum supported length is 64 characters. |
| 23724142 | "records" field must not be specified when a single user is added. |
| 23724143 | SVM UUID, UNIX group name and skip-name-validation are invalid fields for "records" parameter. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

records

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_group

| Name | Type | Description |
|---|---|---|
| name | string | UNIX group name. |

unix_group_users

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

| Name | Type | Description |
|---|---|---|
| records | array[records] | An array of UNIX users specified to add multiple users to a UNIX group in a single API call.<br>Not allowed when the `name` property is used. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX user names is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_group | unix_group | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete a user from a UNIX group**

DELETE `/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users/{name}`

**Introduced In:** 9.9

Deletes a user from the specified UNIX group.

#### Related ONTAP commands

- `vserver services name-service unix-group deluser`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_group.name | string | path | True | UNIX group name. |
| name | string | path | True | UNIX user to be deleted from the group. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 3276897 | The specified UNIX group does not exist in the SVM. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve users for a UNIX group and SVM

GET `/name-services/unix-groups/{svm.uuid}/{unix_group.name}/users/{name}`

**Introduced In:** 9.10

Retrieves users for the specified UNIX group and SVM.

#### Related ONTAP commands

- `vserver services name-service unix-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| unix_group.name | string | path | True | UNIX group name. |
| name | string | path | True | UNIX user to be retrieved from the group. |
| fields | array[string] | query | False | Specify the fields to return. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| unix_group | unix_group | |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "skip_name_validation": true,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_group": {
    "name": "string"
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

records

| Name | Type | Description |
|---|---|---|
| name | string | UNIX user who belongs to the specified UNIX group and the SVM. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_group

| Name | Type | Description |
|---|---|---|
| name | string | UNIX group name. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage UNIX user configurations

### Manage UNIX user configurations

#### Overview

You can use this API to display local UNIX user information and to control UNIX user configurations.

#### Retrieving UNIX user information

You can use the UNIX user GET endpoint to retrieve all of the local UNIX user configurations for data SVMs.

#### Examples

### Retrieving all of the fields for UNIX user configurations for all SVMs

```
# The API:
/api/name-services/unix-users

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-users?fields=*" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
```

```
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "name": "string",
    "id": 7,
    "primary_gid": 0,
    "full_name": "string"
  },
  {
    "svm": {
      "uuid": "3f479a01-2971-11eb-88e1-0050568eefd4",
      "name": "vs2"
    },
    "name": "nobody",
    "id": 65535,
    "primary_gid": 65535,
    "full_name": ""
  },
  {
    "svm": {
      "uuid": "3f479a01-2971-11eb-88e1-0050568eefd4",
      "name": "vs2"
    },
    "name": "pcuser",
    "id": 65534,
    "primary_gid": 65534,
    "full_name": ""
  },
  {
    "svm": {
      "uuid": "3f479a01-2971-11eb-88e1-0050568eefd4",
      "name": "vs2"
    },
    "name": "root",
    "id": 0,
    "primary_gid": 1,
    "full_name": ""
  }
],
"num_records": 4
}
```

### Retrieving the UNIX user configuration of a specific SVM

```
# The API:
/api/name-services/unix-users

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-
users?svm.uuid=25b363a6-2971-11eb-88e1-0050568eefd4&fields=*" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
      "name": "vs1"
    },
    "name": "user1",
    "id": 1,
    "primary_gid": 1,
    "full_name": "Full User Name for user1"
  }
],
"num_records": 1
}
```

### Retrieving a UNIX user configuration of a specific SVM and user

```
# The API:
/api/name-services/unix-users/{svm.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/unix-users/25b363a6-2971-
11eb-88e1-0050568eefd4/user1" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "25b363a6-2971-11eb-88e1-0050568eefd4",
  "name": "vs1"
},
"name": "user1",
"id": 1,
"primary_gid": 1,
"full_name": "Full User Name for user1"
}
```

#### Creating a UNIX user configuration

The UNIX user POST endpoint creates a UNIX user configuration for the specified SVM.

#### Examples

### Creating a UNIX user configuration with all fields specified

```
# The API:
/api/name-services/unix-users

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/unix-users" -H "accept:
application/json" -d "{ \"full_name\": \"Full user name\", \"id\": 7,
\"name\": \"user2\", \"primary_gid\": 10, \"skip_name_validation\": false,
\"svm\": { \"name\": \"vs1\", \"uuid\": \"25b363a6-2971-11eb-88e1-
0050568eefd4\" }}"
```

### Creating a UNIX user configuration with only mandatory fields specified

```
# The API:
/api/name-services/unix-users

# The call:
curl -X POST "https://<mgmt-ip>/api/name-services/unix-users" -H "accept:
application/json" -d "{ \"id\": 8, \"name\": \"user9\", \"primary_gid\":
10, \"svm\": { \"name\": \"vs1\", \"uuid\": \"25b363a6-2971-11eb-88e1-
0050568eefd4\" }}"
```

---

#### Updating a UNIX user configuration

The UNIX user PATCH endpoint updates the UNIX user configuration for the specified user and SVM. The following example shows a PATCH operation:

```
# The API:
/api/name-services/unix-users/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/unix-users/25b363a6-
2971-11eb-88e1-0050568eefd4/user1" -H "accept: application/json" -d "{
\"full_name\": \"Full name\", \"id\": 7, \"primary_gid\": 10 }"
```

---

#### Deleting a UNIX user configuration

The UNIX user DELETE endpoint deletes the specified UNIX user for the specified SVM. The following example shows a DELETE operation:

```
# The API:
/api/name-services/unix-users/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/unix-users/25b363a6-
2971-11eb-88e1-0050568eefd4/user1" -H "accept: application/json"
```

---

### Retrieve all UNIX users for all SVMs

GET `/name-services/unix-users`

**Introduced In:** 9.9

---

Retrieves all of the UNIX users for all of the SVMs.

#### Important notes

- The current UNIX users count can be retrieved from the num_records field by calling the API with the parameter "return_records=false".

#### Related ONTAP commands

- `vserver services name-service unix-user show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name |
| full_name | string | query | False | Filter by full_name |
| id | integer | query | False | Filter by id |
| primary_gid | integer | query | False | Filter by primary_gid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of UNIX user records. |
| records | array[unix_user] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "full_name": "Full User Name for user1",
      "name": "user1",
      "skip_name_validation": true,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_user

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| name | string | UNIX user name to be added to the local database. |

| Name | Type | Description |
|---|---|---|
| primary_gid | integer | Primary group ID to which the user belongs. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a local UNIX user configuration for an SVM

POST `/name-services/unix-users`

**Introduced In:** 9.9

Creates the local UNIX user configuration for an SVM.

#### Important notes

- The default limit for local UNIX users is 32768.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| name | string | UNIX user name to be added to the local database. |
| primary_gid | integer | Primary group ID to which the user belongs. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX user name is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "full_name": "Full User Name for user1",
  "name": "user1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of UNIX user records. |
| records | array[unix_user] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "full_name": "Full User Name for user1",
      "name": "user1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621706 | The specified SVM UUID is incorrect for the specified SVM name |
| 2621516 | This operation is only supported on a data SVM |
| 23724066 | The specified UNIX user name is too long. Maximum supported length is 64 characters. |
| 3277051 | The specified UNIX user name contains invalid characters. Valid characters are 0-9, A-Z, a-z, ".", "_" and "-". |

| Error Code | Description |
| --- | --- |
| 23724128 | The specified UNIX user full-name contains the invalid character ':'. |
| 23724089 | The specified UNIX user full-name is too long. Maximum supported length is 256 characters. |
| 23724055 | Internal error. Failed to create the UNIX user for the SVM. Verify that the cluster is healthy, then try the command again. |
| 23724090 | Configuring individual entries is not supported because file-only configuration is enabled. |
| 3277025 | Failed to create the UNIX user for the SVM because the system limit has been reached. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_user

| Name | Type | Description |
|---|---|---|
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| name | string | UNIX user name to be added to the local database. |
| primary_gid | integer | Primary group ID to which the user belongs. |
| skip_name_validation | boolean | Indicates whether or not the validation for the specified UNIX user name is disabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a UNIX user configuration for an SVM

DELETE `/name-services/unix-users/{svm.uuid}/{name}`

**Introduced In:** 9.9

Deletes a UNIX user configuration for the specified SVM.

#### Related ONTAP commands

- `vserver services name-service unix-user delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX user name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve UNIX user information for a user and SVM

GET `/name-services/unix-users/{svm.uuid}/{name}`

**Introduced In:** 9.9

Retrieves UNIX user information for the specified user and SVM.

#### Related ONTAP commands

- `vserver services name-service unix-user show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX user name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| name | string | UNIX user name to be added to the local database. |
| primary_gid | integer | Primary group ID to which the user belongs. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "full_name": "Full User Name for user1",
  "name": "user1",
  "skip_name_validation": true,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update UNIX user information for a user and SVM

PATCH `/name-services/unix-users/{svm.uuid}/{name}`

**Introduced In:** 9.9

Updates UNIX user information for the specified user and SVM.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | UNIX user name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| primary_gid | integer | Primary group ID to which the user belongs. |

#### Example request

```
{
  "full_name": "Full User Name for user1",
  "skip_name_validation": true
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 23724128 | The specified Unix user full-name contains invalid character ':' |
| 23724089 | The specified UNIX user full-name is too long. Maximum supported length is 256 characters. |
| 23724055 | Internal error. Failed to modify the UNIX user for the SVM. Verify that the cluster is healthy, then try the command again. |
| 23724090 | Configuring individual entries is not supported because file-only configuration is enabled. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

unix_user

| Name | Type | Description |
|---|---|---|
| full_name | string | User's full name. |
| id | integer | UNIX user ID of the specified user. |
| primary_gid | integer | Primary group ID to which the user belongs. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Networking

## ONTAP REST API Networking endpoints

### Overview

The ONTAP networking APIs enable reporting on networking information, such as IPspaces, interfaces, routes, ports, service policies and broadcast domains. Some can also be used to manage networking.

---

### IPspaces

IPspaces enable you to configure a single ONTAP cluster so that it can be accessed by clients from more than one administratively separate network domain, even if those clients are using the same IP address subnet range. This allows for separation of client traffic for privacy and security.

An IPspace defines a distinct IP address space in which storage virtual machines (SVMs) reside. Ports and IP addresses defined for an IPspace are applicable only within that IPspace. A distinct routing table is maintained for each SVM within an IPspace, so that no cross-SVM or cross-IPspace traffic routing occurs.

---

### Ethernet

#### Broadcast Domains

A broadcast domain is a set of ports which would all receive a broadcast packet that is sent from any of the ports.

By accurately representing the physical network with an ONTAP broadcast domain, ONTAP ensures that IP interfaces are able to migrate to appropriate ports in case of failure. ONTAP also ensures that characteristics, such as MTU, stay matched across all ports of the broadcast domain.

A broadcast domain resides in an IPspace, and can be used by cluster-scoped or SVM-scoped IP interfaces in that IPspace. The scope of the broadcast domain’s uniqueness is the IPspace it is in. You must create as many broadcast domains in an IPspace as there are IP subnets with interfaces in that IPspace.

Ports are mapped to an IPspace by assigning the port’s broadcast domain.

#### Ports

A port is a physical or virtual Ethernet network device. Physical ports may be combined into Link Aggregation Groups (LAGs, or ifgrps), or divided into Virtual LANs (VLANs).

The GET and PATCH APIs are available for all port types. The POST and DELETE APIs are available for "lag" and "vlan" port types.

A given port can host zero or more IP interfaces.

A port exists in a broadcast domain and all ports within the same broadcast domain must have layer 2 network connectivity to one another. If a port within a broadcast domain goes down, any IP interfaces hosted by that port can fail over to other ports in the same broadcast domain.

---

### Fibre Channel

#### Interfaces

Fibre Channel (FC) interfaces are the logical endpoints for Fibre Channel network connections to an SVM. A Fibre Channel interface provides Fibre Channel access to storage within the interface's SVM using either Fibre Channel Protocol (FCP) or Non-Volatile Memory Express over Fibre Channel (NVMe over FC).

The Fibre Channel interface REST API allows you to create, delete, update and discover Fibre Channel interfaces and obtain status information for Fibre Channel interfaces.

A Fibre Channel interface is created on a Fibre Channel port that is located on a cluster node. The Fibre Channel port must be specified to identify the location of the interface for a POST or PATCH that relocates an interface. You can identify the port by either supplying the node and port names or the port UUID.

#### Ports

Fibre Channel ports are the physical ports of Fibre Channel adapters on ONTAP cluster nodes that can be connected to Fibre Channel networks to provide Fibre Channel network connectivity. A Fibre Channel port defines the location of a Fibre Channel interface within the ONTAP cluster.

The Fibre Channel port REST API allows you to discover Fibre Channel ports, obtain status information for Fibre Channel ports, and configure Fibre Channel port properties.

#### Learn More

- *Fibre Channel Logins* found in both the *SAN* and *NVMe* sections. Fibre Channel logins represent connections formed by Fibre Channel initiators that have successfully logged in to ONTAP.

---

### IP

#### Interfaces

An interface (also referred to as a *LIF* in ONTAP documentation) represents a network access point to a node in a cluster. In other words, an interface is essentially an IPv4 or IPv6 address with associated attributes.

IP interfaces are configured on ports to send and receive communications over the network. The port that will host the interface can either be explicitly specified using node and/or port fields or implicitly specified using a

broadcast domain.

The IPspace of the interface is required for cluster-scoped interfaces. The SVM owning the interface is required for SVM-scoped interfaces. These interfaces are in the SVM's IPspace.

The service policy of an interface defines what network services are provided by the interface.

#### Routes

Routes indicate which IPv4 or IPv6 gateway to use to communicate with hosts that are not on the local subnet. Typically, an IP interface (or LIF) can only use a gateway if it has the same address family and is in the LIF's subnet.

It is important that every gateway address belongs to a physical or virtual router that has connectivity to the specified destination network.

SVM-scoped routes can only be used by IP interfaces of the specified SVM. Likewise, cluster-scoped routes can only be used by cluster-scoped IP interfaces in the specified IPspace.

#### Service Policies

Service policies are named groupings that define what services are supported by an IP interface. These include both built-in service policies (for example: default-data-files or default-management) and custom service policies.

Service policies are scoped to either an SVM or IPspace.

---

## Manage broadcast domains

### Manage broadcast domains

#### Overview

A broadcast domain is a collection of Ethernet ports that have layer 2 connectivity. They are used to determine which Ethernet ports can host interfaces of various types. The broadcast domain REST API allows you to retrieve, create, modify, and delete broadcast domains. The broadcast domain APIs do not manage port membership. To add a port to a broadcast domain or to move a port to a different broadcast domain, use PATCH /network/ethernet/ports/<uuid>.</uuid>

#### Retrieving network Ethernet broadcast domain information

The broadcast domains GET API retrieves and displays relevant information pertaining to the broadcast domains configured in the cluster. The API retrieves the list of all broadcast domains configured in the cluster, or a specific broadcast domain.

---

#### Examples

### Retrieving all broadcast domains in the cluster

The following output shows the list of all broadcast domains configured in a cluster.

---

```
# The API:
/api/network/ethernet/broadcast-domains

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/broadcast-domains" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "6970c2a9-f34f-11e8-8373-005056bb6b85",
    "name": "Cluster",
    "ipspace": {
      "uuid": "6267eff8-f34f-11e8-8373-005056bb6b85",
      "name": "Cluster",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/6267eff8-f34f-11e8-8373-
005056bb6b85"
        }
      }
    },
    "ports": [
      {
        "uuid": "626b4d19-f34f-11e8-8373-005056bb6b85",
        "name": "e0a",
        "node": {
          "name": "examplecluster-node01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/626b4d19-f34f-11e8-8373-
005056bb6b85"
          }
        }
      },
      {
        "uuid": "626b77b9-f34f-11e8-8373-005056bb6b85",
        "name": "e0b",
        "node": {
          "name": "examplecluster-node01"
        },
        "_links": {
          "self": {
```

```
            "href": "/api/network/ethernet/ports/626b77b9-f34f-11e8-8373-
005056bb6b85"
          }
        }
      }
    ],
    "mtu": 9000,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/6970c2a9-f34f-
11e8-8373-005056bb6b85"
      }
    }
  },
  {
    "uuid": "6972416c-f34f-11e8-8373-005056bb6b85",
    "name": "Default",
    "ipspace": {
      "uuid": "5f650349-f34f-11e8-8373-005056bb6b85",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/5f650349-f34f-11e8-8373-
005056bb6b85"
        }
      }
    },
    "ports": [
      {
        "uuid": "626bae19-f34f-11e8-8373-005056bb6b85",
        "name": "e0c",
        "node": {
          "name": "examplecluster-node01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/626bae19-f34f-11e8-8373-
005056bb6b85"
          }
        }
      },
      {
        "uuid": "626bd677-f34f-11e8-8373-005056bb6b85",
        "name": "e0d",
        "node": {
          "name": "examplecluster-node01"
```

```
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/626bd677-f34f-11e8-8373-
005056bb6b85"
          }
        }
      }
    ],
    "mtu": 1500,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/6972416c-f34f-
11e8-8373-005056bb6b85"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains?fields=*"
  }
}
}
```


#### Retrieving a specific broadcast domain

The following output shows the response returned when a specific broadcast domain is requested. The system returns an error if there is no broadcast domain with the requested UUID.


```
# The API:
/api/network/ethernet/broadcast-domains/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/4475a2c8-f8a0-11e8-8d33-005056bb986f/?fields=*" -H "accept:
application/hal+json"

# The response:
{
"uuid": "4475a2c8-f8a0-11e8-8d33-005056bb986f",
"name": "Cluster",
```

```
"ipspace": {
  "uuid": "3e518ed5-f8a0-11e8-8d33-005056bb986f",
  "name": "Cluster",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/3e518ed5-f8a0-11e8-8d33-005056bb986f"
    }
  }
},
"ports": [
  {
    "uuid": "3e539a62-f8a0-11e8-8d33-005056bb986f",
    "name": "e0a",
    "node": {
      "name": "examplecluster-node01"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/3e539a62-f8a0-11e8-8d33-
005056bb986f"
      }
    }
  },
  {
    "uuid": "3e53c94a-f8a0-11e8-8d33-005056bb986f",
    "name": "e0b",
    "node": {
      "name": "examplecluster-node01"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/3e53c94a-f8a0-11e8-8d33-
005056bb986f"
      }
    }
  }
],
"mtu": 9000,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains/4475a2c8-f8a0-11e8-
8d33-005056bb986f/"
  }
}
}
```

#### Retrieving all broadcast domains with a specific name

The following output shows the response returned when broadcast domains with a specific name in any IPspace are requested.

```
# The API:
/api/network/ethernet/broadcast-domains

# The call:
curl -X GET "https://10.224.87.121/api/network/ethernet/broadcast-
domains/?name=bd1" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "66b607e5-4bee-11e9-af6a-005056bb13c0",
    "name": "bd1",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/66b607e5-4bee-
11e9-af6a-005056bb13c0"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains/?name=bd1"
  }
}
}
```

#### Retrieving the broadcast domains for an IPspace

The following output shows the response returned when the broadcast domains for a specified IPspace are requested.

```
# The API:
/api/network/ethernet/broadcast-domains
```

```
# The call:
curl -X GET "https://10.224.87.121/api/network/ethernet/broadcast-
domains/?ipspace.name=Cluster&fields=*" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "ae69070c-4bed-11e9-af6a-005056bb13c0",
    "name": "Cluster",
    "ipspace": {
      "uuid": "ac466a88-4bed-11e9-af6a-005056bb13c0",
      "name": "Cluster",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/ac466a88-4bed-11e9-af6a-
005056bb13c0"
        }
      }
    },
    "ports": [
      {
        "uuid": "acd67884-4bed-11e9-af6a-005056bb13c0",
        "name": "e0a",
        "node": {
          "name": "examplecluster-node-1"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/acd67884-4bed-11e9-af6a-
005056bb13c0"
          }
        }
      },
      {
        "uuid": "ace1a36f-4bed-11e9-af6a-005056bb13c0",
        "name": "e0b",
        "node": {
          "name": "examplecluster-node-1"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/ace1a36f-4bed-11e9-af6a-
005056bb13c0"
          }
```

```
        }
      }
    ],
    "mtu": 1500,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/ae69070c-4bed-
11e9-af6a-005056bb13c0"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-
domains/?ipspace.name=Cluster&fields=*"
  }
}
}
```


#### Creating network Ethernet broadcast domains

You can use the POST API to create broadcast domains.

#### Example

### Creating a new broadcast domain

The following example shows how to create a broadcast domain with a name of 'bd1' and an MTU of 1500.

```
# The API:
/api/network/ethernet/broadcast-domains

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains?return_records=true" -H "accept: application/hal+json" -d '{
"name": "bd1", "mtu": 1500 }'

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "bd1",
    "mtu": 1500,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/"
      }
    }
  }
]
}
```

#### Updating network Ethernet broadcast domains

You can use the PATCH API to update the attributes of broadcast domains.

#### Examples

### Updating the name and MTU of a specific broadcast domain

The following example shows how the PATCH request changes the broadcast domain name to 'bd2' and the broadcast domain MTU to 9000.

```
# The API:
/api/network/ethernet/broadcast-domains/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/6cde03b2-f8a2-11e8-8d33-005056bb986f/" -d '{ "name": "bd2", "mtu":
9000 }'
{
}
```

#### Updating the IPspace of a specific broadcast domain

The following example shows how the PATCH request changes the IPspace of a broadcast domain to 'ipspace2'.

```
# The API:
/api/network/ethernet/broadcast-domains/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/c6fe2541-61f4-11e9-a66e-005056bbe83e" -d '{ "ipspace" : { "name" :
"ipspace2" } }'
{
}
```

#### Deleting network Ethernet broadcast domains

You can use the DELETE API to delete a broadcast domain from the cluster configuration.

##### Example

#### Deleting a specific broadcast domain

The following DELETE request deletes a broadcast domain.

```
# The API:
/api/network/ethernet/broadcast-domains/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/6cde03b2-f8a2-11e8-8d33-005056bb986f/"
```

---

### Retrieve broadcast domains for the entire cluster

GET `/network/ethernet/broadcast-domains`

**Introduced In:** 9.6

Retrieves a collection of broadcast domains for the entire cluster.

#### Related ONTAP commands

- `network port broadcast-domain show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | query | False | Filter by uuid |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| ports.node.name | string | query | False | Filter by ports.node.name |
| ports.uuid | string | query | False | Filter by ports.uuid |
| ports.name | string | query | False | Filter by ports.name |
| name | string | query | False | Filter by name |
| mtu | integer | query | False | Filter by mtu<br>• Min value: 68 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Min value: 0<br>• Default value: 15<br>• Max value: 120 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[broadcast_domain] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "mtu": 1500,
      "name": "bd1",
      "ports": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
```

```
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

ports

Port UUID along with readable names

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

broadcast_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |
| mtu | integer | Maximum transmission unit, largest packet size on this network |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |
| uuid | string | Broadcast domain UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new broadcast domain**

POST `/network/ethernet/broadcast-domains`

**Introduced In:** 9.6

Creates a new broadcast domain.

#### Required properties

- `name` - Name of the broadcast-domain to create.
- `mtu` - Maximum transmission unit (MTU) of the broadcast domain.

#### Recommended optional properties

- `ipspace.name` or `ipspace.uuid` - IPspace the broadcast domain belongs to.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `ipspace` - *Default*

#### Related ONTAP commands

- `network port broadcast-domain create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |
| mtu | integer | Maximum transmission unit, largest packet size on this network |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |

| Name | Type | Description |
|---|---|---|
| uuid | string | Broadcast domain UUID |

### Example request

```
{
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": 1500,
  "name": "bd1",
  "ports": [
    {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1377267 | The specified IPspace does not exist. |
| 1377596 | The Cluster IPspace cannot contain more than one broadcast domain. Modifying the system IPspace Cluster is not supported. |
| 1966460 | The provided MTU was either too large or too small. |
| 1967082 | The specified ipspace.name does not match the IPspace name of ipspace.uuid. |
| 1967102 | A POST operation might have left the configuration in an inconsistent state. Check the configuration. |
| 1967104 | Invalid IPspace UUID. |
| 53281982 | The specified broadcast domain name is reserved by the system. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

ports

Port UUID along with readable names

| Name | Type | Description |
|---|---|---|
| name | string | |
| node | node | |
| uuid | string | |

broadcast_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |
| mtu | integer | Maximum transmission unit, largest packet size on this network |

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |
| uuid | string | Broadcast domain UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a broadcast domain

DELETE `/network/ethernet/broadcast-domains/{uuid}`

**Introduced In:** 9.6

Deletes a broadcast domain.

#### Related ONTAP commands

- `network port broadcast-domain delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Broadcast domain UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1967103 | A broadcast domain with ports cannot be deleted. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve broadcast domain details

GET `/network/ethernet/broadcast-domains/{uuid}`

**Introduced In:** 9.6

Retrieves details of a broadcast domain.

#### Related ONTAP commands

- `network port broadcast-domain show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Broadcast domain UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |
| mtu | integer | Maximum transmission unit, largest packet size on this network |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |
| uuid | string | Broadcast domain UUID |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": 1500,
  "name": "bd1",
  "ports": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

ports

Port UUID along with readable names

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update broadcast domain properties

PATCH `/network/ethernet/broadcast-domains/{uuid}`

**Introduced In:** 9.6

Updates the properties of a broadcast domain.

#### Related ONTAP commands

- `network port broadcast-domain modify`
- `network port broadcast-domain rename`
- `network port broadcast-domain move`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Broadcast domain UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |

| Name | Type | Description |
|---|---|---|
| mtu | integer | Maximum transmission unit, largest packet size on this network |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |
| uuid | string | Broadcast domain UUID |

**Example request**

```
{
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": 1500,
  "name": "bd1",
  "ports": [
    {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1376484 | Cannot change the MTU of a VLAN to be greater than the MTU of the port hosting it. |
| 1377267 | The specified IPspace does not exist. |
| 1377269 | Failed to lookup the specified IPspace. |
| 1377560 | Broadcast domain already exists in specified IPspace. |
| 1377575 | Remove associated subnets before deleting this broadcast domain. |
| 1377605 | Moving the system-generated broadcast domain to another IPspace is not supported. |
| 1377609 | Updates are partially complete. Updating broadcast domain attributes on some or all of the ports in the broadcast domain have failed. |
| 1966460 | The specified MTU is either too large or too small. |
| 1967082 | The specified ipspace.name does not match the IPspace name of ipspace.uuid. |
| 1967150 | The specified ipspace.uuid is not valid. |
| 1967151 | The specified ipspace.uuid and ipspace.name do not match. |
| 1967152 | Patching IPspace for a broadcast domain requires an effective cluster version of 9.7 or later. |
| 53280884 | The MTU of the broadcast domain cannot be modified on this platform. |
| 53282013 | Broadcast domain cannot be renamed because the name is reserved by the system. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

ports

Port UUID along with readable names

| Name | Type | Description |
|---|---|---|
| name | string | |
| node | node | |
| uuid | string | |

broadcast_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. |
| mtu | integer | Maximum transmission unit, largest packet size on this network |

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| ports | array[ports] | Ports that belong to the broadcast domain |
| uuid | string | Broadcast domain UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network Ethernet ports

### Manage network Ethernet ports

#### Overview

A port is a physical or virtual Ethernet network device. Physical ports may be combined into Link Aggregation Groups (LAGs or ifgrps), or divided into Virtual LANs (VLANs).

GET (collection), GET (instance), and PATCH APIs are available for all port types. POST and DELETE APIs are available for "lag" (ifgrp) and "vlan" port types.

#### Retrieving network port information

The network ports GET API retrieves and displays relevant information pertaining to the ports configured in the cluster. The API retrieves the list of all ports configured in the cluster, or specifically requested ports. The fields returned in the response vary for different ports and configurations.

#### Examples

### Retrieving all ports in the cluster

The following output displays the UUID, name, and port type for all ports configured in a 2-node cluster. The port types are physical, vlan, lag (ifgrp), and p-vlan (available in select environments only).

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ethernet/ports?fields=uuid,name,type" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "2d2c90c0-f70d-11e8-b145-005056bb5b8e",
    "name": "e0a",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/2d2c90c0-f70d-11e8-b145-
005056bb5b8e"
      }
    }
  },
  {
    "uuid": "2d3004da-f70d-11e8-b145-005056bb5b8e",
    "name": "e0b",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/2d3004da-f70d-11e8-b145-
005056bb5b8e"
      }
    }
  },
  {
    "uuid": "2d34a2cb-f70d-11e8-b145-005056bb5b8e",
    "name": "e0c",
    "type": "physical",
    "_links": {
```

```
      "self": {
        "href": "/api/network/ethernet/ports/2d34a2cb-f70d-11e8-b145-
005056bb5b8e"
      }
    }
  },
  {
    "uuid": "2d37189f-f70d-11e8-b145-005056bb5b8e",
    "name": "e0d",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/2d37189f-f70d-11e8-b145-
005056bb5b8e"
      }
    }
  },
  {
    "uuid": "35de5d8b-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0a",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35de5d8b-f70d-11e8-abdf-
005056bb7fc8"
      }
    }
  },
  {
    "uuid": "35de78cc-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0b",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35de78cc-f70d-11e8-abdf-
005056bb7fc8"
      }
    }
  },
  {
    "uuid": "35dead3c-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0c",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35dead3c-f70d-11e8-abdf-
```

```
005056bb7fc8"
      }
    }
  },
  {
    "uuid": "35deda90-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0d",
    "type": "physical",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35deda90-f70d-11e8-abdf-
005056bb7fc8"
      }
    }
  },
  {
    "uuid": "42e25145-f97d-11e8-ade9-005056bb7fc8",
    "name": "e0c-100",
    "type": "vlan",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/42e25145-f97d-11e8-ade9-
005056bb7fc8"
      }
    }
  },
  {
    "uuid": "569e0abd-f97d-11e8-ade9-005056bb7fc8",
    "name": "a0a",
    "type": "lag",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/569e0abd-f97d-11e8-ade9-
005056bb7fc8"
      }
    }
  }
],
"num_records": 10,
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports?fields=uuid,name,type"
  }
}
}
```

### Retrieving a specific physical port

The following output displays the response when a specific physical port is requested. The system returns an error when there is no port with the requested UUID. Also, the "speed" field for the physical port is set only if the state of the port is up.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/2d37189f-f70d-
11e8-b145-005056bb5b8e?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "2d37189f-f70d-11e8-b145-005056bb5b8e",
"name": "e0d",
"mac_address": "00:50:56:bb:62:2d",
"type": "physical",
"node": {
  "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
  "name": "user-cluster-01",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-005056bb5b8e"
    }
  }
},
"broadcast_domain": {
  "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
  "name": "Default",
  "ipspace": {
    "name": "Default"
  },
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-
b145-005056bb5b8e"
    }
  }
},
"enabled": true,
"state": "up",
"mtu": 1500,
"speed": 1000,
```

```
"reachability": "not_repairable",
"reachable_broadcast_domains": [
  {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-
11e8-b145-005056bb5b8e"
      }
    }
  },
  {
    "uuid": "df640ccf-72c4-11ea-b31d-005056bbfb29",
    "name": "Default-1",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/df640ccf-72c4-
11ea-b31d-005056bbfb29"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports/2d37189f-f70d-11e8-b145-
005056bb5b8e"
  }
}
}
```


#### Retrieving a specific VLAN port

The following output displays the response when a specific VLAN port is requested. The system returns an error when there is no port with the requested UUID. Also, the "speed" field for a VLAN port is always set to zero if the state of the port is up. If the state of the port is down, the "speed" field is unset and not reported back.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/42e25145-f97d-
11e8-ade9-005056bb7fc8?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "42e25145-f97d-11e8-ade9-005056bb7fc8",
"name": "e0e-100",
"mac_address": "00:50:56:bb:52:2f",
"type": "vlan",
"node": {
  "uuid": "6042cf47-f70c-11e8-abdf-005056bb7fc8",
  "name": "user-cluster-02",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/6042cf47-f70c-11e8-abdf-005056bb7fc8"
    }
  }
},
"broadcast_domain": {
  "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
  "name": "Default",
  "ipspace": {
    "name": "Default"
  },
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-
b145-005056bb5b8e"
    }
  }
},
"enabled": true,
"state": "up",
"mtu": 1500,
"speed": 0,
"reachability": "ok",
"reachable_broadcast_domains": [
  {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
```

```
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-
11e8-b145-005056bb5b8e"
      }
    }
  }
],
"vlan": {
  "tag": 100,
  "base_port": {
    "uuid": "35deff03-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0e",
    "node": {
      "name": "user-cluster-02"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35deff03-f70d-11e8-abdf-
005056bb7fc8"
      }
    }
  }
},
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports/42e25145-f97d-11e8-ade9-
005056bb7fc8"
  }
}
}
```


---

#### Retrieving a specific LAG port

The following output displays the response when a specific LAG port is requested. The system returns an error when there is no port with the requested UUID. The "lag.active_ports" field is set only if the state of the port is up. Also, the "speed" field for a LAG port is always set to zero if the state of the port is up. If the state of the port is down, the "speed" field is unset and not reported back.

---

```
# The API:
/api/network/ethernet/ports/{uuid}
```

---

```
# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/569e0abd-f97d-
11e8-ade9-005056bb7fc8?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "569e0abd-f97d-11e8-ade9-005056bb7fc8",
"name": "a0a",
"mac_address": "02:50:56:bb:7f:c8",
"type": "lag",
"node": {
  "uuid": "6042cf47-f70c-11e8-abdf-005056bb7fc8",
  "name": "user-cluster-02",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/6042cf47-f70c-11e8-abdf-005056bb7fc8"
    }
  }
},
"broadcast_domain": {
  "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
  "name": "Default",
  "ipspace": {
    "name": "Default"
  },
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-
b145-005056bb5b8e"
    }
  }
},
"enabled": true,
"state": "up",
"mtu": 1500,
"speed": 0,
"reachability": "repairable",
"reachable_broadcast_domains": [
  {
    "uuid": "c7934b4f-691f-11ea-87fd-005056bb1ad3",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
```

```
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/c7934b4f-691f-
11ea-87fd-005056bb1ad3"
      }
    }
  }
],
"lag": {
  "mode": "singlemode",
  "distribution_policy": "mac",
  "member_ports": [
    {
      "uuid": "35df318d-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0f",
      "node": {
        "name": "user-cluster-02"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/35df318d-f70d-11e8-abdf-
005056bb7fc8"
        }
      }
    },
    {
      "uuid": "35df5bad-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0g",
      "node": {
        "name": "user-cluster-02"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/35df5bad-f70d-11e8-abdf-
005056bb7fc8"
        }
      }
    },
    {
      "uuid": "35df9926-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0h",
      "node": {
        "name": "user-cluster-02"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/35df9926-f70d-11e8-abdf-
```

```
005056bb7fc8"
        }
      }
    }
  ],
  "active_ports": [
    {
      "uuid": "35df318d-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0f",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/35df318d-f70d-11e8-abdf-
005056bb7fc8"
        }
      }
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports/569e0abd-f97d-11e8-ade9-
005056bb7fc8"
  }
}
}
```

#### Retrieving all LAG (ifgrp) ports in the cluster

This command retrieves all LAG ports in the cluster (that is, all ports with type=LAG). The example shows how to filter a GET collection based on type.

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ethernet/ports?type=lag&node.name=user-cluster-
01&fields=name,enabled,speed,mtu" -H  "accept: application/hal+json"

# The response:
{
"records": [
  {
```

```
    "uuid": "0c226db0-4b63-11e9-8113-005056bbe040",
    "name": "a0b",
    "type": "lag",
    "node": {
      "name": "user-cluster-01"
    },
    "enabled": true,
    "mtu": 1500,
    "speed": 0,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/0c226db0-4b63-11e9-8113-
005056bbe040"
      }
    }
  },
  {
    "uuid": "d3a84153-4b3f-11e9-a00d-005056bbe040",
    "name": "a0a",
    "type": "lag",
    "node": {
      "name": "user-cluster-01"
    },
    "enabled": true,
    "mtu": 1500,
    "speed": 0,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/d3a84153-4b3f-11e9-a00d-
005056bbe040"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/network/ethernet/ports?fields=name,enabled,speed,mtu&type=lag&node.n
ame=user-cluster-01"
  }
}
}
```

#### Creating VLAN and LAG ports

You can use the network ports POST API to create VLAN and LAG ports. If you supply the optional broadcast domain property, the specified broadcast domain will be assigned to the new port immediately. Otherwise, within a few minutes automatic probing will determine the correct broadcast domain and will assign it to the port. During that period of time, the port will not be capable of hosting interfaces.

---

#### Examples

### Creating a VLAN port

The following output displays the record returned after the creation of a VLAN port on "e0e" and VLAN tag "100".

---

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"type\":
\"vlan\", \"node\": { \"name\": \"user-cluster-01\" }, \"enabled\": true,
\"vlan\": { \"tag\": 100, \"base_port\": { \"name\": \"e0e\", \"node\": {
\"name\": \"user-cluster-01\" } } }}"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "88b2f682-fa42-11e8-a6d7-005056bb5b8e",
    "type": "vlan",
    "node": {
      "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
      "name": "user-cluster-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-
005056bb5b8e"
        }
      }
    },
    "enabled": true,
    "vlan": {
      "tag": 100,
```

```
      "base_port": {
        "uuid": "2d39df72-f70d-11e8-b145-005056bb5b8e",
        "name": "e0e",
        "node": {
          "name": "user-cluster-01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/2d39df72-f70d-11e8-b145-
005056bb5b8e"
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/88b2f682-fa42-11e8-a6d7-
005056bb5b8e"
      }
    }
  }
]
}
```


#### Creating a VLAN port in a specific broadcast domain

The following output displays the record returned after the creation of a VLAN port on "e0e" and VLAN tag "100". Also, the VLAN port is added to the "Default" broadcast domain in the "Default" IPspace.

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"type\": \"vlan\", \"node\": { \"name\": \"user-cluster-01\" },
\"broadcast_domain\": { \"name\": \"Default\", \"ipspace\": { \"name\":
\"Default    \" } }, \"enabled\": true, \"vlan\": { \"tag\": 100,
\"base_port\": { \"name\": \"e0e\", \"node\": { \"name\": \"user-cluster-
01\" } } }}"

# The response:
```

```
{
"num_records": 1,
"records": [
  {
    "uuid": "88b2f682-fa42-11e8-a6d7-005056bb5b8e",
    "type": "vlan",
    "node": {
      "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
      "name": "user-cluster-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-
005056bb5b8e"
        }
      }
    },
    "broadcast_domain": {
      "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
      "name": "Default",
      "ipspace": {
        "name": "Default"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-
11e8-b145-005056bb5b8e"
        }
      }
    },
    "enabled": true,
    "vlan": {
      "tag": 100,
      "base_port": {
        "uuid": "2d39df72-f70d-11e8-b145-005056bb5b8e",
        "name": "e0e",
        "node": {
          "name": "user-cluster-01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/2d39df72-f70d-11e8-b145-
005056bb5b8e"
          }
        }
      }
    },
```

```
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/88b2f682-fa42-11e8-a6d7-
005056bb5b8e"
      }
    }
  }
]
}
```

#### Creating a LAG (ifgrp) port

The following output displays the record returned after the creation of a LAG port with "e0f", "e0g" and "e0h" as member ports.

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{  \"type\":
\"lag\",  \"node\": { \"name\": \"user-cluster-01\"  }, \"enabled\": true,
\"lag\": { \"mode\": \"singlemode\", \"distribution_policy\": \"mac\",
\"member_ports\": [ { \"name\": \"e0f\", \"node\": { \"name\": \"user-
cluster-01\" } }, { \"name\": \"e0g\", \"node\": { \"name\": \"user-
cluster-01\" }}, { \"name\": \"e0h\", \"node\": { \"name\": \"user-
cluster-01\" } } ]  } }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "1807772a-fa4d-11e8-a6d7-005056bb5b8e",
    "type": "lag",
    "node": {
      "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
      "name": "user-cluster-01"
    },
    "enabled": true,
    "lag": {
      "mode": "singlemode",
```

```
      "distribution_policy": "mac",
      "member_ports": [
        {
          "uuid": "2d3c9adc-f70d-11e8-b145-005056bb5b8e",
          "name": "e0f",
          "node": {
            "name": "user-cluster-01"
          }
        },
        {
          "uuid": "2d40b097-f70d-11e8-b145-005056bb5b8e",
          "name": "e0g",
          "node": {
            "name": "user-cluster-01"
          }
        },
        {
          "uuid": "2d46d01e-f70d-11e8-b145-005056bb5b8e",
          "name": "e0h",
          "node": {
            "name": "user-cluster-01"
          }
        }
      ]
    }
  }
]
}
```


#### Creating a LAG (ifgrp) port in a specific broadcast domain

The following output displays the record returned after the creation of a LAG port with "e0f", "e0g" and "e0h" as member ports. Also, the LAG port is added to the "Default" broadcast domain in the "Default" IPspace.

```
# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{  \"type\":
\"lag\",  \"node\": { \"name\": \"user-cluster-01\"  },
\"broadcast_domain\": { \"name\": \"Default\", \"ipspace\": { \"name\":
```

```
\"Default\" } }, \"enabled\": true, \"lag\": { \"mode\": \"singlemode\",
\"distribution_policy\": \"mac\", \"member_ports\": [ { \"name\": \"e0f\",
\"node\": { \"name\": \"user-cluster-01\" } }, { \"name\": \"e0g\",
\"node\": { \"name\": \"user-cluster-01\" }}, { \"name\": \"e0h\",
\"node\": { \"name\": \"user-cluster-01\" } } ]  } }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "1807772a-fa4d-11e8-a6d7-005056bb5b8e",
    "type": "lag",
    "node": {
      "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
      "name": "user-cluster-01"
    },
    "broadcast_domain": {
      "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
      "name": "Default",
      "ipspace": {
        "name": "Default"
      }
    },
    "enabled": true,
    "lag": {
      "mode": "singlemode",
      "distribution_policy": "mac",
      "member_ports": [
        {
          "uuid": "2d3c9adc-f70d-11e8-b145-005056bb5b8e",
          "name": "e0f",
          "node": {
            "name": "user-cluster-01"
          }
        },
        {
          "uuid": "2d40b097-f70d-11e8-b145-005056bb5b8e",
          "name": "e0g",
          "node": {
            "name": "user-cluster-01"
          }
        },
        {
          "uuid": "2d46d01e-f70d-11e8-b145-005056bb5b8e",
          "name": "e0h",
```

```
            "node": {
              "name": "user-cluster-01"
            }
          }
        ]
      }
    }
  ]
  }
```

#### Updating ports

You can use the network ports PATCH API to update the attributes of ports.

#### Examples

### Updating the broadcast domain of a port

The following PATCH request removes the port from the current broadcast domain and adds it to the specified broadcast domain.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/6867efaf-d702-
11e8-994f-005056bbc994" -H  "accept: application/hal+json" -H  "Content-
Type: application/json" -d "{ \"broadcast_domain\": { \"name\":
\"Default\", \"ipspace\": { \"name\": \"Default\" }}}"
```

### Updating the admin status of a port

The following PATCH request brings the specified port down.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/51d3ab39-d86d-
11e8-aca6-005056bbc994" -H  "accept: application/hal+json" -H  "Content-
Type: application/json" -d "{ \"enabled\": \"false\" }"
```

### Repairing a port

The following PATCH request repairs a port. Only ports that have reachability as "repairable" can be repaired. The "reachability" parameter cannot be patched in the same request as other parameters that might affect the target port's reachability status.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/51d3ab39-d86d-
11e8-aca6-005056bbc994" -H  "accept: application/hal+json" -H  "Content-
Type: application/json" -d "{ \"reachability\": \"ok\" }"
```

#### Deleting ports

You can use the network ports DELETE API to delete VLAN and LAG ports in the cluster. Note that physical ports cannot be deleted.
Deleting a port also removes the port from the broadcast domain.

#### Example

### Deleting a VLAN port

The network ports DELETE API is used to delete a VLAN port.

```
# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ethernet/ports/6867efaf-
d702-11e8-994f-005056bbc994" -H  "accept: application/hal+json" -H
"Content-Type: application/json"
```

---

### Retrieve ports

GET `/network/ethernet/ports`

**Introduced In:** 9.6

Retrieves a collection of ports (physical, VLAN and LAG) for an entire cluster.

#### Related ONTAP commands

- `network port show`
- `network port ifgrp show`
- `network port vlan show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| state | string | query | False | Filter by state |
| name | string | query | False | Filter by name |
| mac_address | string | query | False | Filter by mac_address |
| speed | integer | query | False | Filter by speed |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.device.timestamp | string | query | False | Filter by statistics.device.timestamp<br><br>• Introduced in: 9.8 |
| statistics.device.receive_raw.discards | integer | query | False | Filter by statistics.device.receive_raw.discards<br><br>• Introduced in: 9.8 |
| statistics.device.receive_raw.errors | integer | query | False | Filter by statistics.device.receive_raw.errors<br><br>• Introduced in: 9.8 |
| statistics.device.receive_raw.packets | integer | query | False | Filter by statistics.device.receive_raw.packets<br><br>• Introduced in: 9.8 |
| statistics.device.transmit_raw.discards | integer | query | False | Filter by statistics.device.transmit_raw.discards<br><br>• Introduced in: 9.8 |
| statistics.device.transmit_raw.errors | integer | query | False | Filter by statistics.device.transmit_raw.errors<br><br>• Introduced in: 9.8 |
| statistics.device.transmit_raw.packets | integer | query | False | Filter by statistics.device.transmit_raw.packets<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.device.link_down_count_raw | integer | query | False | Filter by statistics.device.link_down_count_raw<br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br>• Introduced in: 9.8 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.8 |
| lag.active_ports.node.name | string | query | False | Filter by lag.active_ports.node.name |
| lag.active_ports.uuid | string | query | False | Filter by lag.active_ports.uuid |
| lag.active_ports.name | string | query | False | Filter by lag.active_ports.name |
| lag.distribution_policy | string | query | False | Filter by lag.distribution_policy |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lag.mode | string | query | False | Filter by lag.mode |
| lag.member_ports.node.name | string | query | False | Filter by lag.member_ports.node.name |
| lag.member_ports.uuid | string | query | False | Filter by lag.member_ports.uuid |
| lag.member_ports.name | string | query | False | Filter by lag.member_ports.name |
| type | string | query | False | Filter by type |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| reachability | string | query | False | Filter by reachability<br>• Introduced in: 9.8 |
| enabled | boolean | query | False | Filter by enabled |
| interface_count | integer | query | False | Filter by interface_count<br>• Introduced in: 9.11 |
| vlan.base_port.node.name | string | query | False | Filter by vlan.base_port.node.name |
| vlan.base_port.uuid | string | query | False | Filter by vlan.base_port.uuid |
| vlan.base_port.name | string | query | False | Filter by vlan.base_port.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| vlan.tag | integer | query | False | Filter by vlan.tag<br>• Max value: 4094<br>• Min value: 1 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br>• Introduced in: 9.8 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br>• Introduced in: 9.8 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.8 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.8 |
| uuid | string | query | False | Filter by uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| broadcast_domain.name | string | query | False | Filter by broadcast_domain.name |
| broadcast_domain.uuid | string | query | False | Filter by broadcast_domain.uuid |
| broadcast_domain.ipspace.name | string | query | False | Filter by broadcast_domain.ipspace.name |
| mtu | integer | query | False | Filter by mtu<br>• Min value: 68 |
| discovered_devices.remaining_hold_time | integer | query | False | Filter by discovered_devices.remaining_hold_time<br>• Introduced in: 9.11 |
| discovered_devices.protocol | string | query | False | Filter by discovered_devices.protocol<br>• Introduced in: 9.11 |
| discovered_devices.version | string | query | False | Filter by discovered_devices.version<br>• Introduced in: 9.11 |
| discovered_devices.ip_addresses | string | query | False | Filter by discovered_devices.ip_addresses<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| discovered_devices.chassis_id | string | query | False | Filter by discovered_devices.chassis_id<br><br>• Introduced in: 9.11 |
| discovered_devices.capabilities | string | query | False | Filter by discovered_devices.capabilities<br><br>• Introduced in: 9.11 |
| discovered_devices.platform | string | query | False | Filter by discovered_devices.platform<br><br>• Introduced in: 9.11 |
| discovered_devices.remote_port | string | query | False | Filter by discovered_devices.remote_port<br><br>• Introduced in: 9.11 |
| discovered_devices.system_name | string | query | False | Filter by discovered_devices.system_name<br><br>• Introduced in: 9.11 |
| discovered_devices.name | string | query | False | Filter by discovered_devices.name<br><br>• Introduced in: 9.11 |
| reachable_broadcast_domains.name | string | query | False | Filter by reachable_broadcast_domains.name<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| reachable_broadcast_domains.uuid | string | query | False | Filter by reachable_broadcast_domains.uuid<br>• Introduced in: 9.8 |
| reachable_broadcast_domains.ipspace.name | string | query | False | Filter by reachable_broadcast_domains.ipspace.name<br>• Introduced in: 9.8 |
| pfc_queues_admin | integer | query | False | Filter by pfc_queues_admin<br>• Introduced in: 9.16<br>• Max value: 7<br>• Min value: 0 |
| flowcontrol_admin | string | query | False | Filter by flowcontrol_admin<br>• Introduced in: 9.16 |
| rdma_protocols | string | query | False | Filter by rdma_protocols<br>• Introduced in: 9.10 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[port] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ipspace": {
          "name": "ipspace1"
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "discovered_devices": [
        {
          "capabilities": [
            "router",
            "switch"
          ],
          "chassis_id": "string",
          "ip_addresses": [
            "192.168.100.24",
            "192.168.100.26"
          ],
          "name": "ETY-R1S4-510Q13.datacenter.example.com",
          "platform": "93180YC-EX",
          "protocol": "cdp",
          "remote_port": "FastEthernet0/12",
```

```
          "system_name": "string",
          "version": "Cisco Nexus Operating System (NX-OS) Software,
Version 8.1"
        }
      ],
      "flowcontrol_admin": "string",
      "interface_count": 0,
      "lag": {
        "active_ports": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "distribution_policy": "string",
        "member_ports": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "mode": "string"
      },
      "mac_address": "01:02:03:04:05:06",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
        "duration": "PT15S",
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "mtu": 1500,
      "name": "e1b",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "pfc_queues_admin": [
        "integer"
      ],
      "rdma_protocols": [
        "roce"
      ],
      "reachability": "ok",
      "reachable_broadcast_domains": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ipspace": {
            "name": "ipspace1"
          },
          "name": "bd1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "speed": 1000,
      "state": "string",
      "statistics": {
        "device": {
          "link_down_count_raw": 3,
```

```
          "receive_raw": {
            "discards": 100,
            "errors": 200,
            "packets": 500
          },
          "timestamp": "2017-01-25 11:20:13 +0000",
          "transmit_raw": {
            "discards": 100,
            "errors": 200,
            "packets": 500
          }
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "type": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vlan": {
        "base_port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "tag": 100
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain's IPspace |

broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

discovered_devices

| Name | Type | Description |
|---|---|---|
| capabilities | array[string] | The list of the capabilities of the discovered device. |

| Name | Type | Description |
| --- | --- | --- |
| chassis_id | string | Identifier associated with this specific discovered device, useful for locating the device in a data center. |
| ip_addresses | array[string] | The IP addresses on the discovered device. |
| name | string | Name of the discovered device. |
| platform | string | Hardware platform of the discovered device. |
| protocol | string | The protocol used to identify the discovered device. This can have a value of CDP or LLDP. |
| remaining_hold_time | integer | The number of seconds until the discovered device entry expires and is removed. |
| remote_port | string | The name of the remote port on the discovered device. The format is dependent on the reporting device. |
| system_name | string | Additional name used to identify a specific piece of equipment. |
| version | string | The version of the software running on the discovered device. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

active_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| node | node | |
| uuid | string | |

member_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

lag

| Name | Type | Description |
| --- | --- | --- |
| active_ports | array[active_ports] | Active ports of a LAG (ifgrp). (Some member ports may be inactive.) |
| distribution_policy | string | Policy for mapping flows to ports for outbound packets through a LAG (ifgrp). |
| member_ports | array[member_ports] | Array of ports belonging to the LAG, regardless of their state. |
| mode | string | Determines how the ports interact with the switch. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the port.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

reachable_broadcast_domains

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

device

Device-related counters for the port object. These counters are applicable at the lowest layer of the

networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

| Name | Type | Description |
|---|---|---|
| link_down_count_raw | integer | The number of link state changes from up to down seen on the device. |
| receive_raw | receive_raw | Packet receive counters for the Ethernet port. |
| timestamp | string | The timestamp when the device specific counters were collected. |
| transmit_raw | transmit_raw | Packet transmit counters for the Ethernet port. |

throughput_raw

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the port.

| Name | Type | Description |
| --- | --- | --- |
| device | device | Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the throughput_raw performance data. |

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

vlan

| Name | Type | Description |
|---|---|---|
| base_port | base_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| tag | integer | VLAN ID |

port

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| flowcontrol_admin | string | Requested flow control |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |
| mac_address | string | |
| metric | metric | The most recent sample of I/O metrics for the port. |

| Name | Type | Description |
| --- | --- | --- |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| node | node | |
| pfc_queues_admin | array[integer] | List of PFC queues |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| reachability | string | Reachability status of the port. Enum value "ok" is the only acceptable value for a PATCH request to repair a port. |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| statistics | statistics | The real time I/O statistics for the port. |
| type | string | Type of physical or virtual port |
| uuid | string | Port UUID |
| vlan | vlan | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a new VLAN or LAG

POST `/network/ethernet/ports`

**Introduced In:** 9.6

Creates a new VLAN (such as node1:e0a-100) or LAG (ifgrp, such as node2:a0a).

#### Required properties

- `node` - Node the port will be created on.
- `vlan` - This field cannot be specified at the same time as `lag`.
  - `vlan.base_port` - Physical port or LAG the VLAN will be created on.
  - `vlan.tag` - Tag used to identify VLAN on the base port.
- `lag` - This field cannot be specified at the same time as `vlan`.
  - `lag.mode` - Policy for the LAG that will be created.
  - `lag.distribution_policy` - Indicates how the packets are distributed between ports.
  - `lag.member_ports` - Set of ports the LAG consists of.

#### Optional properties

- `type` - Defines if a VLAN or LAG will be created:
- `broadcast_domain` - The layer-2 broadcast domain the port is associated with. The port will be placed in a broadcast domain if it is not specified. It may take several minutes for the broadcast domain to be assigned. During that period the port cannot host interfaces.

#### Related ONTAP commands

- `network port ifgrp create`
- `network port vlan create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |
| mac_address | string | |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| node | node | |
| rdma_protocols | array[string] | Supported RDMA offload protocols |

| Name | Type | Description |
| --- | --- | --- |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| type | string | Type of physical or virtual port |
| uuid | string | Port UUID |
| vlan | vlan | |

**Example request**


```
{
  "broadcast_domain": {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "discovered_devices": [
    {
      "capabilities": [
        "router",
        "switch"
      ],
      "chassis_id": "string",
      "ip_addresses": [
        "192.168.100.24",
        "192.168.100.26"
      ],
      "name": "ETY-R1S4-510Q13.datacenter.example.com",
      "platform": "93180YC-EX",
      "protocol": "cdp",
      "remote_port": "FastEthernet0/12",
      "system_name": "string",
      "version": "Cisco Nexus Operating System (NX-OS) Software,
Version 8.1"
    }
  ],
  "interface_count": 0,
  "lag": {
    "active_ports": [
      {
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "distribution_policy": "string",
    "member_ports": [
      {
        "name": "e1b",
        "node": {
```

```
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "mode": "string"
    },
    "mac_address": "01:02:03:04:05:06",
    "mtu": 1500,
    "name": "e1b",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "rdma_protocols": [
      "roce"
    ],
    "reachable_broadcast_domains": [
      {
        "ipspace": {
          "name": "ipspace1"
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "speed": 1000,
    "state": "string",
    "type": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "vlan": {
      "base_port": {
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "tag": 100
    }
  }
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[port] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "broadcast_domain": {
        "ipspace": {
          "name": "ipspace1"
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "discovered_devices": [
        {
          "capabilities": [
            "router",
            "switch"
          ],
          "chassis_id": "string",
          "ip_addresses": [
            "192.168.100.24",
            "192.168.100.26"
          ],
          "name": "ETY-R1S4-510Q13.datacenter.example.com",
          "platform": "93180YC-EX",
          "protocol": "cdp",
          "remote_port": "FastEthernet0/12",
          "system_name": "string",
          "version": "Cisco Nexus Operating System (NX-OS) Software,
Version 8.1"
        }
      ],
      "interface_count": 0,
      "lag": {
        "active_ports": [
          {
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "distribution_policy": "string",
        "member_ports": [
```

```
            {
              "name": "e1b",
              "node": {
                "name": "node1"
              },
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          ],
          "mode": "string"
        },
        "mac_address": "01:02:03:04:05:06",
        "mtu": 1500,
        "name": "e1b",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "rdma_protocols": [
          "roce"
        ],
        "reachable_broadcast_domains": [
          {
            "ipspace": {
              "name": "ipspace1"
            },
            "name": "bd1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "speed": 1000,
        "state": "string",
        "type": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
        "vlan": {
          "base_port": {
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "tag": 100
        }
      }
    ]
  }
```

**Headers**

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1376361 | Port is already a member of a LAG. |
| 1376857 | Cannot create a VLAN on a port that is a member of a LAG. |
| 1377609 | The update is partially complete. Updating broadcast domain attributes on this port has failed. |
| 1966189 | Port is the home port or current port of an interface. |
| 1966466 | VLAN ID must be a number from 1 to 4094. |
| 1967083 | The specified type is not valid. |
| 1967084 | The specified node UUID is not valid. |
| 1967085 | The specified node name is not valid. |
| 1967086 | Node name and UUID must match if both are provided. |
| 1967087 | The specified broadcast domain UUID is not valid. |
| 1967088 | The specified broadcast domain name does not exist in the specified IPspace. |
| 1967089 | The specified broadcast domain UUID, name, and IPspace name do not match. |
| 1967090 | The specified VLAN base port UUID is not valid. |
| 1967091 | The specified VLAN base port name and node name are not valid. |
| 1967092 | The specified node does not match the node specified for the VLAN base port. |
| 1967093 | The specified VLAN base port UUID, name, and VLAN base port node name do not match. |
| 1967094 | The specified LAG member port UUID is not valid. |
| 1967095 | The specified LAG member port name and node name combination is not valid. |

| Error Code | Description |
| --- | --- |
| 1967096 | The specified node does not match the specified LAG member port node. |
| 1967097 | The specified LAG member ports UUID, name, and node name do not match. |
| 1967098 | VLAN POST operation has failed because admin status could not be set for the specified port. |
| 1967099 | Partial success of the VLAN POST operation. Verify the state of the created VLAN for more information. |
| 1967100 | LAG POST operation failed because admin status could not be set. |
| 1967101 | Partial success of the LAG POST operation. Verify the state of the created LAG for more information. |
| 1967102 | POST operation might have left the configuration in an inconsistent state. Check the configuration. |
| 1967126 | A LAG requires at least one member port. |
| 1967148 | Failure to remove port from broadcast domain. |
| 1967149 | Failure to add port to broadcast domain. |
| 1967175 | VLANs cannot be created on ports in the Cluster IPspace. |
| 1967190 | Missing or incomplete VLAN specification. |
| 1967191 | Missing or incomplete lag specification. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain's IPspace |

broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

discovered_devices

| Name | Type | Description |
|---|---|---|
| capabilities | array[string] | The list of the capabilities of the discovered device. |
| chassis_id | string | Identifier associated with this specific discovered device, useful for locating the device in a data center. |
| ip_addresses | array[string] | The IP addresses on the discovered device. |
| name | string | Name of the discovered device. |
| platform | string | Hardware platform of the discovered device. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocol used to identify the discovered device. This can have a value of CDP or LLDP. |
| remaining_hold_time | integer | The number of seconds until the discovered device entry expires and is removed. |
| remote_port | string | The name of the remote port on the discovered device. The format is dependent on the reporting device. |
| system_name | string | Additional name used to identify a specific piece of equipment. |
| version | string | The version of the software running on the discovered device. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

active_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

member_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

lag

| Name | Type | Description |
| --- | --- | --- |
| active_ports | array[active_ports] | Active ports of a LAG (ifgrp). (Some member ports may be inactive.) |
| distribution_policy | string | Policy for mapping flows to ports for outbound packets through a LAG (ifgrp). |
| member_ports | array[member_ports] | Array of ports belonging to the LAG, regardless of their state. |
| mode | string | Determines how the ports interact with the switch. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the port.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

reachable_broadcast_domains

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Broadcast domain UUID |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

device

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

| Name | Type | Description |
| --- | --- | --- |
| link_down_count_raw | integer | The number of link state changes from up to down seen on the device. |
| timestamp | string | The timestamp when the device specific counters were collected. |

throughput_raw

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate

of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the port.

| Name | Type | Description |
|---|---|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the throughput_raw performance data. |

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

vlan

| Name | Type | Description |
| --- | --- | --- |
| base_port | base_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| tag | integer | VLAN ID |

port

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |
| mac_address | string | |

| Name | Type | Description |
| --- | --- | --- |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| node | node | |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| type | string | Type of physical or virtual port |
| uuid | string | Port UUID |
| vlan | vlan | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a VLAN or LAG

DELETE `/network/ethernet/ports/{uuid}`

**Introduced In:** 9.6

Deletes a VLAN or LAG.

#### Related ONTAP commands

- `network port ifgrp delete`
- `network port vlan delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Port UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1376858 | Port already has an interface bound. |
| 1966189 | Port is the home port or current port of an interface. |
| 1966302 | This interface group is hosting VLAN interfaces that must be deleted before running this command. |
| 1967105 | Cannot delete a physical port. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a physical port, VLAN, or LAG details

GET `/network/ethernet/ports/{uuid}`

**Introduced In:** 9.6

Retrieves the details of a physical port, VLAN, or LAG.

#### Related ONTAP commands

- `network port show`
- `network port ifgrp show`
- `network port vlan show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Port UUID |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| flowcontrol_admin | string | Requested flow control |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |
| mac_address | string | |
| metric | metric | The most recent sample of I/O metrics for the port. |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| node | node | |
| pfc_queues_admin | array[integer] | List of PFC queues |
| rdma_protocols | array[string] | Supported RDMA offload protocols |

| Name | Type | Description |
| --- | --- | --- |
| reachability | string | Reachability status of the port. Enum value "ok" is the only acceptable value for a PATCH request to repair a port. |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| statistics | statistics | The real time I/O statistics for the port. |
| type | string | Type of physical or virtual port |
| uuid | string | Port UUID |
| vlan | vlan | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "broadcast_domain": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "discovered_devices": [
    {
      "capabilities": [
        "router",
        "switch"
      ],
      "chassis_id": "string",
      "ip_addresses": [
        "192.168.100.24",
        "192.168.100.26"
      ],
      "name": "ETY-R1S4-510Q13.datacenter.example.com",
      "platform": "93180YC-EX",
      "protocol": "cdp",
      "remote_port": "FastEthernet0/12",
      "system_name": "string",
      "version": "Cisco Nexus Operating System (NX-OS) Software,
Version 8.1"
    }
  ],
  "flowcontrol_admin": "string",
  "interface_count": 0,
  "lag": {
    "active_ports": [
      {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "distribution_policy": "string",
    "member_ports": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "mode": "string"
  },
  "mac_address": "01:02:03:04:05:06",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "mtu": 1500,
  "name": "e1b",
```

```
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "pfc_queues_admin": [
    "integer"
  ],
  "rdma_protocols": [
    "roce"
  ],
  "reachability": "ok",
  "reachable_broadcast_domains": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipspace": {
        "name": "ipspace1"
      },
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "speed": 1000,
  "state": "string",
  "statistics": {
    "device": {
      "link_down_count_raw": 3,
      "receive_raw": {
        "discards": 100,
        "errors": 200,
        "packets": 500
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "transmit_raw": {
        "discards": 100,
        "errors": 200,
        "packets": 500
      }
```

```
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "type": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "vlan": {
      "base_port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "tag": 100
    }
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain's IPspace |

broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

discovered_devices

| Name | Type | Description |
|---|---|---|
| capabilities | array[string] | The list of the capabilities of the discovered device. |
| chassis_id | string | Identifier associated with this specific discovered device, useful for locating the device in a data center. |
| ip_addresses | array[string] | The IP addresses on the discovered device. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the discovered device. |
| platform | string | Hardware platform of the discovered device. |
| protocol | string | The protocol used to identify the discovered device. This can have a value of CDP or LLDP. |
| remaining_hold_time | integer | The number of seconds until the discovered device entry expires and is removed. |
| remote_port | string | The name of the remote port on the discovered device. The format is dependent on the reporting device. |
| system_name | string | Additional name used to identify a specific piece of equipment. |
| version | string | The version of the software running on the discovered device. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

active_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

member_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

lag

| Name | Type | Description |
|---|---|---|
| active_ports | array[active_ports] | Active ports of a LAG (ifgrp). (Some member ports may be inactive.) |
| distribution_policy | string | Policy for mapping flows to ports for outbound packets through a LAG (ifgrp). |
| member_ports | array[member_ports] | Array of ports belonging to the LAG, regardless of their state. |
| mode | string | Determines how the ports interact with the switch. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the port.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

reachable_broadcast_domains

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
|---|---|---|
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
|---|---|---|
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

device

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

| Name | Type | Description |
| --- | --- | --- |
| link_down_count_raw | integer | The number of link state changes from up to down seen on the device. |
| receive_raw | receive_raw | Packet receive counters for the Ethernet port. |
| timestamp | string | The timestamp when the device specific counters were collected. |
| transmit_raw | transmit_raw | Packet transmit counters for the Ethernet port. |

throughput_raw

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the port.

| Name | Type | Description |
| --- | --- | --- |
| device | device | Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the throughput_raw performance data. |

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

vlan

| Name | Type | Description |
| --- | --- | --- |
| base_port | base_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| tag | integer | VLAN ID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a port

PATCH `/network/ethernet/ports/{uuid}`

**Introduced In:** 9.6

Updates a port.

#### Related ONTAP commands

- `network port broadcast-domain add-ports`
- `network port broadcast-domain remove-ports`
- `network port modify`
- `network port ifgrp add-port`
- `network port ifgrp remove-port`
- `network port reachability repair`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Port UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| flowcontrol_admin | string | Requested flow control |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |
| mac_address | string | |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| pfc_queues_admin | array[integer] | List of PFC queues |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| reachability | string | Reachability status of the port. Enum value "ok" is the only acceptable value for a PATCH request to repair a port. |

| Name | Type | Description |
| --- | --- | --- |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| uuid | string | Port UUID |

**Example request**


```
{
  "broadcast_domain": {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "discovered_devices": [
    {
      "capabilities": [
        "router",
        "switch"
      ],
      "chassis_id": "string",
      "ip_addresses": [
        "192.168.100.24",
        "192.168.100.26"
      ],
      "name": "ETY-R1S4-510Q13.datacenter.example.com",
      "platform": "93180YC-EX",
      "protocol": "cdp",
      "remote_port": "FastEthernet0/12",
      "system_name": "string",
      "version": "Cisco Nexus Operating System (NX-OS) Software,
Version 8.1"
    }
  ],
  "flowcontrol_admin": "string",
  "interface_count": 0,
  "lag": {
    "active_ports": [
      {
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "member_ports": [
      {
        "name": "e1b",
        "node": {
```

```
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ]
    },
    "mac_address": "01:02:03:04:05:06",
    "mtu": 1500,
    "name": "e1b",
    "pfc_queues_admin": [
      "integer"
    ],
    "rdma_protocols": [
      "roce"
    ],
    "reachability": "ok",
    "reachable_broadcast_domains": [
      {
        "ipspace": {
          "name": "ipspace1"
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "speed": 1000,
    "state": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1376361 | Port is already a member of a LAG. |

| Error Code | Description |
| --- | --- |
| 1376369 | Cannot add port to the interface group because the port hosts one or more VLANs. |
| 1376488 | Disabling the last operational cluster port on a node is not allowed. |
| 1376492 | Cannot change the MTU of a VLAN to be greater than the MTU of the port hosting it. |
| 1377563 | Port is already a member of a LAG. |
| 1377608 | Port cannot be used because it is currently the home port or current port of an interface. |
| 1377609 | The update is partially complete. Updating broadcast domain attributes on this port has failed. |
| 1966288 | Disabling the cluster ports can only be done on the local node. |
| 1967087 | The specified broadcast domain UUID is not valid. |
| 1967088 | The specified broadcast domain name does not exist in the specified IPspace. |
| 1967089 | The specified broadcast domain UUID, name and IPspace name do not match. |
| 1967094 | The specified LAG member port UUID is not valid. |
| 1967095 | The specified LAG member port name and node name combination is not valid. |
| 1967096 | The specified node does not match the specified LAG member port node. |
| 1967097 | The specified LAG member ports UUID, name, and node name do not match. |
| 1967126 | A LAG requires at least one member port. |
| 1967148 | Failure to remove port from broadcast domain. |
| 1967149 | Failure to add port to broadcast domain. |
| 1967184 | The reachability parameter cannot be patched in the same request as other parameters that might affect the target port's reachability status. |
| 1967185 | The port cannot be repaired because the port is deemed as non-repairable. |
| 1967186 | Invalid value for the reachability parameter. |
| 1967195 | Missing or incomplete name retrieval for specified port UUID. |
| 1967580 | This command is not supported as the effective cluster version is earlier than 9.8. |

| Error Code | Description |
| --- | --- |
| 1967582 | The reachability parameter is not supported on this cluster. |
| 53216932 | Failed to determine whether newly introduced PFC flow control is supported. |
| 53280899 | Not all nodes support enabling the PFC feature. |
| 53280900 | No PFC queues specified. |
| 53280901 | Failed to modify PFC. |
| 53280902 | Cluster ports does not support PFC flow control. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain's IPspace |

broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

discovered_devices

| Name | Type | Description |
|---|---|---|
| capabilities | array[string] | The list of the capabilities of the discovered device. |
| chassis_id | string | Identifier associated with this specific discovered device, useful for locating the device in a data center. |
| ip_addresses | array[string] | The IP addresses on the discovered device. |
| name | string | Name of the discovered device. |
| platform | string | Hardware platform of the discovered device. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocol used to identify the discovered device. This can have a value of CDP or LLDP. |
| remaining_hold_time | integer | The number of seconds until the discovered device entry expires and is removed. |
| remote_port | string | The name of the remote port on the discovered device. The format is dependent on the reporting device. |
| system_name | string | Additional name used to identify a specific piece of equipment. |
| version | string | The version of the software running on the discovered device. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

active_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

member_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

lag

| Name | Type | Description |
| --- | --- | --- |
| active_ports | array[active_ports] | Active ports of a LAG (ifgrp). (Some member ports may be inactive.) |
| member_ports | array[member_ports] | Array of ports belonging to the LAG, regardless of their state. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the port.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

reachable_broadcast_domains

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | |
| name | string | Name of the broadcast domain, scoped to its IPspace |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Broadcast domain UUID |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

device

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

| Name | Type | Description |
| --- | --- | --- |
| link_down_count_raw | integer | The number of link state changes from up to down seen on the device. |
| timestamp | string | The timestamp when the device specific counters were collected. |

throughput_raw

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate

of throughput bytes per unit of time.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the port.

| Name | Type | Description |
|---|---|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the throughput_raw performance data. |

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

vlan

| Name | Type | Description |
| --- | --- | --- |
| base_port | base_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| tag | integer | VLAN ID |

port

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input. |
| discovered_devices | array[discovered_devices] | Discovered devices |
| enabled | boolean | |
| flowcontrol_admin | string | Requested flow control |
| interface_count | integer | Number of interfaces hosted. This field is only applicable for cluster administrators. No value is returned for SVM administrators. If the node hosting a port is not healthy no value will be returned. |
| lag | lag | |

| Name | Type | Description |
| --- | --- | --- |
| mac_address | string | |
| mtu | integer | MTU of the port in bytes. Set by broadcast domain. |
| name | string | Portname, such as e0a, e1b-100 (VLAN on Ethernet), a0c (LAG/ifgrp), a0d-200 (VLAN on LAG/ifgrp), e0a.pv1 (p-VLAN, in select environments only) |
| pfc_queues_admin | array[integer] | List of PFC queues |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| reachability | string | Reachability status of the port. Enum value "ok" is the only acceptable value for a PATCH request to repair a port. |
| reachable_broadcast_domains | array[reachable_broadcast_domains] | Reachable broadcast domains. |
| speed | integer | Link speed in Mbps |
| state | string | Operational state of the port. The state is set to 'down' if the operational state of the port is down. The state is set to 'up' if the link state of the port is up and the port is healthy. The state is set to 'up' if the link state of the port is up and configured to ignore health status. The state is 'degraded' if the link state of the port is up, and the port is not healthy. |
| uuid | string | Port UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical port performance metrics**

GET `/network/ethernet/ports/{uuid}/metrics`

**Introduced In:** 9.8

Retrieves historical performance metrics for a port.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| timestamp | string | query | False | Filter by timestamp |
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |
| uuid | string | path | True | Unique identifier of the port. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Throughput performance for the Ethernet port.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | Port UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve network Ethernet switch ports

### Retrieve network Ethernet switch ports

#### Overview

This API can be used to get the port information for an ethernet switch used in a cluster or storage networks. This API supports GET only. The GET operation returns a list of ports with status and configuration information.

#### Examples

### Retrieving the ports for ethernet switches

The following example retrieves the ethernet switch ports for all the ethernet switches used for cluster and/or storage networks.
Note that if the *fields=** parameter is not specified, the fields identity.number, statistics, and mac_address are not returned.
Filters can be added on the fields to limit the results.

```
# The API:
GET /network/ethernet/switch/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switch/ports?fields=*"
-H  "accept: application/hal+json" -H  "Content-Type:
application/hal+json"

# The response: (abbreviated output due to length, shows a port connected
to a node, a port with no remote connection and a port connected to shelf)
{
"records": [
  {
    "switch": {
      "name": "RTP-CS01-510R11(FOC22092K12)",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/switches/RTP-CS01-
510R11(FOC22092K12)"
```

```
        }
      }
    },
    "identity": {
      "name": "Ethernet1/1",
      "index": 436207616,
      "number": 1,
      "type": "ethernetcsmacd"
    },
    "mtu": 9216,
    "duplex_type": "full_duplex",
    "speed": 100000,
    "configured": "up",
    "state": "up",
    "isl": false,
    "vpc_peer_link": false,
    "statistics": {
      "timestamp": "2024-11-08T19:51:12+00:00",
      "receive_raw": {
        "packets": 1616467751,
        "errors": 0,
        "discards": 0
      },
      "transmit_raw": {
        "packets": 206717534,
        "errors": 0,
        "discards": 0
      }
    },
    "remote_port": {
      "mtu": 9000,
      "name": "e3a",
      "functional_roles": [
        "cluster",
        "ha"
      ],
      "device": {
        "discovered_name": "stiA400-311",
        "node": {
          "name": "stiA400-311",
          "uuid": "54c0f036-8a3a-11ea-893d-00a098fd726d",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/54c0f036-8a3a-11ea-893d-
00a098fd726d"
            }
```

```
          }
        }
      }
    },
    "mac_address": "00:be:75:ae:2a:d4",
    "vlan_id": [
      1,
      17,
      18
    ],
    "roles": [
      {
        "type": "cluster",
        "zone": 1
      }
    ],
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switch/ports/RTP-CS01-
510R11%28FOC22092K12%29/Ethernet1%2F1/436207616"
      }
    }
  },
  {
    "switch": {
      "name": "RTP-CS01-510R11(FOC22092K12)",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/switches/RTP-CS01-
510R11(FOC22092K12)"
        }
      }
    },
    "identity": {
      "name": "Ethernet1/11",
      "index": 436212736,
      "number": 11,
      "type": "ethernetcsmacd"
    },
    "mtu": 9216,
    "duplex_type": "unknown",
    "speed": 100000,
    "configured": "up",
    "state": "down",
    "isl": false,
    "vpc_peer_link": false,
```

```
    "statistics": {
      "timestamp": "2024-11-08T19:51:12+00:00",
      "receive_raw": {
        "packets": 0,
        "errors": 0,
        "discards": 0
        },
      "transmit_raw": {
        "packets": 0,
        "errors": 0,
        "discards": 0
      }
    },
    "mac_address": "00be75ae2afc",
    "vlan_id": [
      1,
      17,
      18
    ],
    "roles": [
      {
        "type": "cluster",
        "zone": 1
      }
    ],
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switch/ports/RTP-CS01-
510R11%28FOC22092K12%29/Ethernet1%2F11/436212736"
      }
    }
  },
  {
    "switch": {
      "name": "RTP-SS01-510R10(FOC22170DFR)",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/switches/RTP-SS01-
510R10(FOC22170DFR)"
        }
      }
    },
    "identity": {
      "name": "Ethernet1/10",
      "index": 436212224,
      "number": 10,
```

```
      "type": "ethernetcsmacd"
    },
    "mtu": 9216,
    "duplex_type": "full_duplex",
    "speed": 100000,
    "configured": "up",
    "state": "up",
    "isl": false,
    "vpc_peer_link": false,
    "statistics": {
      "timestamp": "2024-11-08T19:51:12+00:00",
      "receive_raw": {
        "packets": 332013844,
        "errors": 0,
        "discards": 0
      },
      "transmit_raw": {
        "packets": 2429595607,
        "errors": 0,
        "discards": 0
      }
    },
    "remote_port": {
      "mtu": 9000,
      "name": "e0a",
      "functional_roles": [
        "storage_shelf"
      ],
      "device": {
        "discovered_name": "1.1",
        "shelf": {
          "name": "1.1",
          "module": "B",
          "uid": "12439000444923584512",
          "_links": {
            "self": {
              "href": "/api/storage/shelves/12439000444923584512"
            }
          }
        }
      }
    },
    "mac_address": "00fcbaead548",
    "vlan_id": [
      1,
      30
```

```
    ],
    "roles": [
      {
        "type": "storage",
        "zone": 1
      }
    ],
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switch/ports/RTP-SS01-
510R10%28FOC22170DFR%29/Ethernet1%2F10/436212224"
      }
    }
  }
],
"num_records": 138,
"_links": {
  "self": {
    "href": "/api/network/ethernet/switch/ports?fields=*"
  }
}
}
```


---

#### Retrieving a ports on an ethernet switch

```
# The API:
GET /network/ethernet/switch/ports/{switch.name}/{identity.name}/
{identity.index}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switch/ports/RTP-SS02-
510R10%28FOC22131U6T%29/Ethernet1%2F9/436211712" -H  "accept:
application/hal+json" -H  "Content-Type: application/hal+json"

# The response:
{
"switch": {
  "name": "RTP-SS02-510R10(FOC22131U6T)",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switches/RTP-SS02-
510R10(FOC22131U6T)"
    }
  }
```

```
},
"identity": {
  "name": "Ethernet1/9",
  "index": 436211712,
  "number": 9,
  "type": "ethernetcsmacd"
},
"mtu": 9216,
"duplex_type": "full_duplex",
"speed": 100000,
"configured": "up",
"state": "up",
"isl": false,
"vpc_peer_link": false,
"statistics": {
  "timestamp": "2024-11-08T19:51:12+00:00",
  "receive_raw": {
    "packets": 4012559315,
    "errors": 0,
    "discards": 0
  },
  "transmit_raw": {
    "packets": 337898026,
    "errors": 0,
    "discards": 0
 }
},
"remote_port": {
  "mtu": 9000,
  "name": "e0a",
  "functional_roles": [
    "storage_shelf"
  ],
  "device": {

    # cSpell:disable-next-line
    "discovered_name": "SHFFG1234567890:A",
    "shelf": {
      "name": "1.1",
      "module": "A",
      "uid": "12439000444923584512",
      "_links": {
        "self": {
          "href": "/api/storage/shelves/12439000444923584512"
        }
      }
```

```
      }
    }
  },
  "mac_address": "00fcbaea7228",
  "vlan_id": [
    1,
    30
  ],
  "roles": [
    {
      "type": "storage",
      "zone": 1
    }
  ],
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switch/ports/RTP-SS02-
510R10%28FOC22131U6T%29/Ethernet1%2F9/436211712"
    }
  }
}
```


---

### Retrieve Ethernet switch ports

GET `/network/ethernet/switch/ports`

**Introduced In:** 9.8

Retrieves the ethernet switch ports.

#### Related ONTAP commands

- `system switch ethernet interface show`

#### Learn more

- DOC /network/ethernet/switch/ports

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| switch.name | string | query | False | Filter by switch.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| identity.breakout.physical_port | string | query | False | Filter by identity.breakout.physical_port<br>• Introduced in: 9.17 |
| identity.breakout.number | integer | query | False | Filter by identity.breakout.number<br>• Introduced in: 9.17 |
| identity.number | integer | query | False | Filter by identity.number |
| identity.index | integer | query | False | Filter by identity.index |
| identity.name | string | query | False | Filter by identity.name |
| vpc_peer_link | boolean | query | False | Filter by vpc_peer_link<br>• Introduced in: 9.15 |
| mtu | integer | query | False | Filter by mtu |
| configured | string | query | False | Filter by configured |
| isl | boolean | query | False | Filter by isl |
| duplex_type | string | query | False | Filter by duplex_type |
| roles.type | string | query | False | Filter by roles.type<br>• Introduced in: 9.17 |
| roles.dr_group | string | query | False | Filter by roles.dr_group<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| roles.zone | integer | query | False | Filter by roles.zone<br>• Introduced in: 9.17<br>• Min value: 1 |
| mac_address | string | query | False | Filter by mac_address |
| speed | integer | query | False | Filter by speed |
| state | string | query | False | Filter by state |
| remote_port.functional_roles | string | query | False | Filter by remote_port.functional_roles<br>• Introduced in: 9.17 |
| remote_port.device.node.name | string | query | False | Filter by remote_port.device.node.name |
| remote_port.device.node.uuid | string | query | False | Filter by remote_port.device.node.uuid |
| remote_port.device.discovered_name | string | query | False | Filter by remote_port.device.discovered_name<br>• Introduced in: 9.17 |
| remote_port.device.dcn.serial_number | string | query | False | Filter by remote_port.device.dcn.serial_number<br>• Introduced in: 9.18 |
| remote_port.device.dcn.name | string | query | False | Filter by remote_port.device.dcn.name<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| remote_port.device.dcn.uuid | string | query | False | Filter by remote_port.device.dcn.uuid<br><br>• Introduced in: 9.18 |
| remote_port.device.shelf.name | string | query | False | Filter by remote_port.device.shelf.name<br><br>• Introduced in: 9.12 |
| remote_port.device.shelf.uid | string | query | False | Filter by remote_port.device.shelf.uid |
| remote_port.device.shelf.module | string | query | False | Filter by remote_port.device.shelf.module<br><br>• Introduced in: 9.12 |
| remote_port.mtu | integer | query | False | Filter by remote_port.mtu |
| remote_port.name | string | query | False | Filter by remote_port.name |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.17 |
| statistics.transmit_raw.discards | integer | query | False | Filter by statistics.transmit_raw.discards |
| statistics.transmit_raw.errors | integer | query | False | Filter by statistics.transmit_raw.errors |
| statistics.transmit_raw.packets | integer | query | False | Filter by statistics.transmit_raw.packets |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.receive_raw.discards | integer | query | False | Filter by statistics.receive_raw.discards |
| statistics.receive_raw.errors | integer | query | False | Filter by statistics.receive_raw.errors |
| statistics.receive_raw.packets | integer | query | False | Filter by statistics.receive_raw.packets |
| type | string | query | False | Filter by type |
| vlan_id | integer | query | False | Filter by vlan_id |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Min value: 0<br>• Default value: 15<br>• Max value: 120 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[switch_port] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "configured": "string",
      "duplex_type": "string",
      "identity": {
        "breakout": {
          "number": 1,
          "physical_port": "Ethernet1/9"
        },
        "index": 0,
        "name": "string",
        "number": 0
      },
      "mac_address": "string",
      "mtu": 0,
      "remote_port": {
        "device": {
          "dcn": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "serial_number": "4048820-60-9",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "discovered_name": "stiA400-311",
          "node": {
```

```
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "shelf": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "module": "string",
            "name": "1.1",
            "uid": "12439000444923584512"
          }
        },
        "functional_roles": [
          "string"
        ],
        "mtu": 0,
        "name": "string"
      },
      "roles": [
        {
          "dr_group": "string",
          "type": "string"
        }
      ],
      "speed": 0,
      "state": "string",
      "statistics": {
        "receive_raw": {
          "discards": 100,
          "errors": 200,
          "packets": 500
        },
        "timestamp": "2024-11-18 15:52:17 +0000",
        "transmit_raw": {
          "discards": 100,
          "errors": 200,
          "packets": 500
        }
      },
```

```
      "switch": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "RTP-SS01-510R03(FOC223443KQ)"
      },
      "type": "string",
      "vlan_id": [
        "integer"
      ]
    }
  ]
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

breakout

| Name | Type | Description |
|---|---|---|
| number | integer | Breakout port sub-interface number. |
| physical_port | string | Breakout physical port name. |

identity

| Name | Type | Description |
|---|---|---|
| breakout | breakout | |
| index | integer | Interface Index. |
| name | string | Interface Name. |
| number | integer | Interface Number. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dcn

Compute node connected to this port.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Compute node name. |
| serial_number | string | Compute node serial number. |
| uuid | string | Compute node UUID. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

shelf

Shelf connected to this port.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| module | string | Shelf module connected to this port. |
| name | string | |
| uid | string | |

device

Device connected to port.

| Name | Type | Description |
| --- | --- | --- |
| dcn | dcn | Compute node connected to this port. |
| discovered_name | string | Raw name of the discovered device. |
| node | node | |
| shelf | shelf | Shelf connected to this port. |

remote_port

Remote port.

| Name | Type | Description |
|---|---|---|
| device | device | Device connected to port. |
| functional_roles | array[string] | |
| mtu | integer | MTU in octets. |
| name | string | Remote port name. |

roles

Allowed use for this port.

| Name | Type | Description |
|---|---|---|
| dr_group | string | DR group. |
| type | string | Allowed use type. |
| zone | integer | Zone ID to differentiate between roles with the same type. |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
|---|---|---|
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
|---|---|---|
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |

| Name | Type | Description |
| --- | --- | --- |
| packets | integer | Total packet count. |

statistics

These are raw counters for the device associated with the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| receive_raw | receive_raw | Packet receive counters for the Ethernet port. |
| timestamp | string | The time the statistics were gathered. |
| transmit_raw | transmit_raw | Packet transmit counters for the Ethernet port. |

switch

The name of the specified cluster or storage switch.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | |

switch_port

Ethernet Switch Port REST API

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| configured | string | Administrative Status. |
| duplex_type | string | Duplex Settings. |
| identity | identity | |
| isl | boolean | Is configured as an ISL link. |
| mac_address | string | MAC Address. |
| mtu | integer | MTU. |
| remote_port | remote_port | Remote port. |

| Name | Type | Description |
|---|---|---|
| roles | array[roles] | |
| speed | integer | Interface Speed(Mbps). |
| state | string | Operational Status. |
| statistics | statistics | These are raw counters for the device associated with the Ethernet port. |
| switch | switch | The name of the specified cluster or storage switch. |
| type | string | Interface Type. |
| vlan_id | array[integer] | |
| vpc_peer_link | boolean | Is configured as a Virtual Port Channel (vPC) peer-link. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve an Ethernet switch port**

GET `/network/ethernet/switch/ports/{switch}/{identity.name}/{identity.index}`

**Introduced In:** 9.8

Retrieves an ethernet switch port.

#### Related ONTAP commands

- `system switch ethernet interface show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| switch | string | path | True | Switch Name |
| identity.name | string | path | True | Interface Name |
| identity.index | integer | path | True | Interface Index |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| configured | string | Administrative Status. |
| duplex_type | string | Duplex Settings. |
| identity | identity | |
| isl | boolean | Is configured as an ISL link. |
| mac_address | string | MAC Address. |
| mtu | integer | MTU. |
| remote_port | remote_port | Remote port. |
| roles | array[roles] | |
| speed | integer | Interface Speed(Mbps). |
| state | string | Operational Status. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw counters for the device associated with the Ethernet port. |
| switch | switch | The name of the specified cluster or storage switch. |
| type | string | Interface Type. |
| vlan_id | array[integer] | |
| vpc_peer_link | boolean | Is configured as a Virtual Port Channel (vPC) peer-link. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configured": "string",
  "duplex_type": "string",
  "identity": {
    "breakout": {
      "number": 1,
      "physical_port": "Ethernet1/9"
    },
    "index": 0,
    "name": "string",
    "number": 0
  },
  "mac_address": "string",
  "mtu": 0,
  "remote_port": {
    "device": {
      "dcn": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "serial_number": "4048820-60-9",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "discovered_name": "stiA400-311",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "shelf": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "module": "string",
        "name": "1.1",
        "uid": "12439000444923584512"
      }
    },
    "functional_roles": [
      "string"
    ],
    "mtu": 0,
    "name": "string"
  },
  "roles": [
    {
      "dr_group": "string",
      "type": "string"
    }
  ],
  "speed": 0,
  "state": "string",
  "statistics": {
    "receive_raw": {
      "discards": 100,
      "errors": 200,
      "packets": 500
    },
    "timestamp": "2024-11-18 15:52:17 +0000",
    "transmit_raw": {
      "discards": 100,
      "errors": 200,
      "packets": 500
    }
  },
  "switch": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "RTP-SS01-510R03(FOC223443KQ)"
  },
  "type": "string",
  "vlan_id": [
    "integer"
```

```
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

breakout

| Name | Type | Description |
|---|---|---|
| number | integer | Breakout port sub-interface number. |
| physical_port | string | Breakout physical port name. |

identity

| Name | Type | Description |
|---|---|---|
| breakout | breakout | |
| index | integer | Interface Index. |
| name | string | Interface Name. |
| number | integer | Interface Number. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

dcn

Compute node connected to this port.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Compute node name. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | Compute node serial number. |
| uuid | string | Compute node UUID. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

shelf

Shelf connected to this port.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| module | string | Shelf module connected to this port. |
| name | string | |
| uid | string | |

device

Device connected to port.

| Name | Type | Description |
| --- | --- | --- |
| dcn | dcn | Compute node connected to this port. |
| discovered_name | string | Raw name of the discovered device. |
| node | node | |
| shelf | shelf | Shelf connected to this port. |

remote_port

Remote port.

| Name | Type | Description |
| --- | --- | --- |
| device | device | Device connected to port. |
| functional_roles | array[string] | |
| mtu | integer | MTU in octets. |
| name | string | Remote port name. |

roles

Allowed use for this port.

| Name | Type | Description |
| --- | --- | --- |
| dr_group | string | DR group. |
| type | string | Allowed use type. |
| zone | integer | Zone ID to differentiate between roles with the same type. |

receive_raw

Packet receive counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

transmit_raw

Packet transmit counters for the Ethernet port.

| Name | Type | Description |
| --- | --- | --- |
| discards | integer | Total number of discarded packets. |
| errors | integer | Number of packet errors. |
| packets | integer | Total packet count. |

statistics

These are raw counters for the device associated with the Ethernet port.

| Name | Type | Description |
|---|---|---|
| receive_raw | receive_raw | Packet receive counters for the Ethernet port. |
| timestamp | string | The time the statistics were gathered. |
| transmit_raw | transmit_raw | Packet transmit counters for the Ethernet port. |

switch

The name of the specified cluster or storage switch.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage Ethernet switches

## Manage Ethernet switches

### Overview

This endpoint manages Ethernet switches used by ONTAP for monitoring and configuration. You can retrieve, configure, add, or delete Ethernet switches, as well as change SNMP credentials and monitor their status.

### Examples

## Retrieving a list of Ethernet switches

The following example shows the response with a list of monitored Ethernet switches. Note that if the *fields=** parameter is not specified, the fields snmp.version, snmp.user, version, monitoring.enabled, and monitoring.reason are not returned.

```
# The API:
/network/ethernet/switches

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switches?fields=*" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json"

# The response:
{
"records": [
  {
    "name": "RTP-CS01-510R11(FOC22092K12)",
    "address": "172.26.207.77",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "enabled": true,
      "monitored": true,
      "reason": "None"
    },
    "network": "cluster",
    "serial_number": "Unknown",
    "snmp": {
      "version": "snmpv2c",
      "user": "cshm1!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.2(3)",
    "role": "cluster",
    "rcf_version": "RCF NX3232C v1.13 1-CLUSTER",
    "_links": {
```

```
      "self": {
        "href": "/api/network/ethernet/switches/RTP-CS01-
510R11%28FOC22092K12%29"
      }
    }
  },
  {
    "name": "RTP-CS01-510R12(FOC22373C3P)",
    "address": "172.26.207.82",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "enabled": true,
      "monitored": true,
      "reason": "None"
    },
    "network": "cluster",
    "serial_number": "FOC22373C3P",
    "snmp": {
      "version": "snmpv2c",
      "user": "cshm1!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.2(3)",
    "role": "cluster",
    "rcf_version": "RCF NX3232C v1.13 1-CLUSTER",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-CS01-
510R12%28FOC22373C3P%29"
      }
    }
  },
  {
    "name": "RTP-SS01-510R10(FOC22170DFR)",
    "address": "172.26.207.65",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "enabled": true,
      "monitored": true,
      "reason": "None"
    },
    "network": "storage",
    "serial_number": "FOC22170DFR",
    "snmp": {
```

```
      "version": "snmpv2c",
      "user": "cshm1!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.3(3)",
    "role": "storage",
    "rcf_version": "RCF NX3232C v1.13 1-STORAGE",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-SS01-
510R10%28FOC22170DFR%29"
      }
    }
  },
  {
    "name": "RTP-SS02-510R10(FOC22131U6T)",
    "address": "172.26.207.66",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "enabled": true,
      "monitored": true,
      "reason": "None"
    },
    "network": "storage",
    "serial_number": "FOC22131U6T",
    "snmp": {
      "version": "snmpv2c",
      "user": "cshm1!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.3(3)",
    "role": "storage",
    "rcf_version": "RCF NX3232C v1.13 1-STORAGE",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-SS02-
510R10%28FOC22131U6T%29"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/network/ethernet/switches?fields=*"
```

```
  }
}
}
```

#### Retrieving an Ethernet switch

The following example shows the response of the requested Ethernet switch. If there is no Ethernet switch with the requested name, an error is returned.

```
# The API:
/network/ethernet/switches/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switches/RTP-SS02-
510R10(FOC22131U6T)?fields=*" -H  "accept: application/hal+json" -H
"Content-Type: application/hal+json"

# The response:
{
"name": "RTP-SS02-510R10(FOC22131U6T)",
"address": "172.26.207.66",
"discovered": true,
"model": "NX3232C",
"monitoring": {
  "enabled": true,
  "monitored": true,
  "reason": "None"
},
"network": "storage",
"serial_number": "FOC22131U6T",
"snmp": {
  "version": "snmpv2c",
  "user": "cshm1!"
},
"version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.3(3)",
"role": "storage",
"rcf_version": "RCF NX3232C v1.13 1-STORAGE",
"_links": {
  "self": {
    "href": "/api/network/ethernet/switches/RTP-SS02-510R10(FOC22131U6T)"
  }
}
}
```

### Configuring monitoring settings for an Ethernet switch

The following example configures SNMP credential and version settings for an Ethernet switch. If the provided information is not valid, an error is returned.

```
# The API:
/network/ethernet/switches/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/switches/sconqa-
corduroyl-03" -H "accept: application/hal+json" -H "Content-Type:
application/hal+json" -d '{"snmp": {"version": "snmpv2c", "user":
"cshm1!"}}'

# The response:
{
"job": {
  "uuid": "f84fbb3c-589c-11ec-b181-d039ea48a7dc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f84fbb3c-589c-11ec-b181-d039ea48a7dc"
    }
  }
}
}
```

### Adding an Ethernet switch

The following example adds an Ethernet switch for monitoring. If the provided information is not valid, an error is returned.

```
# The API:
/network/ethernet/switches

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ethernet/switches" -H "accept:
application/hal+json" -H "Content-Type: application/hal+json" -d '{"name":
"RTP-SS02-510R10(FOC22131U6T)", "address": "172.26.207.66", "model":
"NX3232C", "monitoring": {"enabled": "true"}, "network": "storage",
"snmp": {"version": "snmpv2c", "user": "cshm1!"}}'

# The response:
{
"job": {
  "uuid": "f84fbb3c-589c-11ec-b181-d039ea48a7dc",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f84fbb3c-589c-11ec-b181-d039ea48a7dc"
    }
  }
}
}
```

#### Deleting an Ethernet switch

The following example deletes an Ethernet switch. If there is no Ethernet switch with the requested name, an error is returned.

```
# The API:
/network/ethernet/switches/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ethernet/switches/sconqa-
corduroyl-03" -H "accept: application/hal+json" -H "Content-Type:
application/hal+json"

# The response:
{
}
```

#### Retrieve Ethernet switches attached to a chassis

GET `/network/ethernet/switches`

**Introduced In:** 9.8

Retrieves a collection of Ethernet switches.

#### Related ONTAP commands

- `system switch ethernet show`

#### Learn more

- DOC /network/ethernet/switches

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| monitoring.enabled | boolean | query | False | Filter by monitoring.enabled |
| monitoring.monitored | boolean | query | False | Filter by monitoring.monitored |
| monitoring.reason | string | query | False | Filter by monitoring.reason |
| snmp.version | string | query | False | Filter by snmp.version |
| snmp.user | string | query | False | Filter by snmp.user |
| role | string | query | False | Filter by role<br>• Introduced in: 9.17 |
| serial_number | string | query | False | Filter by serial_number |
| model | string | query | False | Filter by model |
| address | string | query | False | Filter by address |
| name | string | query | False | Filter by name |
| discovered | boolean | query | False | Filter by discovered |
| network | string | query | False | Filter by network |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rcf_version | string | query | False | Filter by rcf_version<br>• Introduced in: 9.17 |
| version | string | query | False | Filter by version |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | collection_links | |
| num_records | integer | Number of Records |
| records | array[switch] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "address": "string",
      "model": "string",
      "monitoring": {
        "reason": "string"
      },
      "name": "string",
      "network": "string",
      "rcf_version": "RCF NX9336C-FX2 v1.13 1-CLUSTER",
      "role": [
        "cluster+storage"
      ],
      "serial_number": "string",
      "snmp": {
        "user": "string",
        "version": "string"
      },
      "version": "string"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

monitoring

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable Health Monitoring. |
| monitored | boolean | Is Monitored. |
| reason | string | Reason For Not Monitoring. |

snmp

| Name | Type | Description |
|---|---|---|
| user | string | Community String or SNMPv3 Username. |
| version | string | SNMP Version. |

switch

Ethernet Switch REST API

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| address | string | IP Address. |

| Name | Type | Description |
| --- | --- | --- |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| model | string | Model Number. |
| monitoring | monitoring | |
| name | string | Name. |
| network | string | Switch Network. |
| rcf_version | string | The switch reference configuration file (RCF) version. |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles. Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an Ethernet switch

POST `/network/ethernet/switches`

**Introduced In:** 9.11

Adds an Ethernet switch.

#### Required properties

- `name` - Name of the switch to create.
- `address` - Switch IP address.
- `model` - Switch model number.
- `monitoring.enabled` - Whether the switch should be monitored by CSHM.
- `network`
  - *cluster* for cluster or shared switches.
  - *storage* for storage switches.
- `snmp.version` - SNMP version.
- `snmp.user` - SNMP user.

#### Related ONTAP commands

- `system switch ethernet create`

#### Learn more

- DOC /network/ethernet/switches

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| address | string | IP Address. |

| Name | Type | Description |
| --- | --- | --- |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| model | string | Model Number. |
| monitoring | monitoring | |
| name | string | Name. |
| network | string | Switch Network. |
| rcf_version | string | The switch reference configuration file (RCF) version. |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles.<br>Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

### Example request

```
{
  "address": "string",
  "model": "string",
  "monitoring": {
    "reason": "string"
  },
  "name": "string",
  "network": "string",
  "rcf_version": "RCF NX9336C-FX2 v1.13 1-CLUSTER",
  "role": [
    "cluster+storage"
  ],
  "serial_number": "string",
  "snmp": {
    "user": "string",
    "version": "string"
  },
  "version": "string"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5636148 | Failed to discover SNMPv3 engine ID of the remote switch. |
| 5636149 | The "<snmp.user>\" must be a valid SNMPv3 user belonging to remote switch \"<name.>\" and must also be configured in ONTAP.</name.></snmp.user> |
| 5636152 | 0.0.0.0 cannot be specified as the IP address for the "<address.>\" parameter.</address.> |
| 8912900 | Invalid IP address. |
| 9044009 | SNMPv3 client feature is not available until the effective cluster version is ONTAP 9.3 or later. |
| 12517377 | Model number validation failed. Specify the correct model number and try the command again. |
| 12517378 | The IP "<address.>\" is not reachable. Verify that the address is valid, or check the network path.</address.> |
| 12517380 | SNMP validation request timed out. Verify that the "<snmp.user>\" parameter is valid.</snmp.user> |
| 12517382 | SNMPv3 validation request timed out. Verify that the "<snmp.user>\" parameter is valid (i.e., the SNMPv3 user exists in ONTAP and on the remote switch). If the \"<snmp.user>\" parameter is valid, verify that the SNMPv3 user’s credentials are the same both in ONTAP as well as in the remote switch. If a custom engine-id was provided for the SNMPv3 user, ensure it is the same as that of the remote switch.</snmp.user></snmp.user> |
| 12517384 | A switch with the same IP address is already being monitored. |

| Error Code | Description |
| --- | --- |
| 12517385 | The "<model.>\" is unknown. Use \"OTHER\" if the switch model is not one of the following: {models}.</model.> |
| 12517411 | The input value for "<name.>\", \"<address.>\" and/or \"<network.>\" must match its discovered value.</network.></address.></name.> |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

monitoring

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable Health Monitoring. |
| monitored | boolean | Is Monitored. |
| reason | string | Reason For Not Monitoring. |

snmp

| Name | Type | Description |
|---|---|---|
| user | string | Community String or SNMPv3 Username. |
| version | string | SNMP Version. |

switch

Ethernet Switch REST API

| Name | Type | Description |
|---|---|---|
| address | string | IP Address. |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| model | string | Model Number. |
| monitoring | monitoring | |
| name | string | Name. |

| Name | Type | Description |
|---|---|---|
| network | string | Switch Network. |
| rcf_version | string | The switch reference configuration file (RCF) version. |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles. Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete an Ethernet switch

DELETE `/network/ethernet/switches/{name}`

**Introduced In:** 9.11

Deletes an Ethernet switch.

#### Related ONTAP commands

- `system switch ethernet delete`

#### Learn more

- DOC /network/ethernet/switches

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | path | True | Switch Name. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 12517412 | Cannot remove a discovered switch. To stop monitoring the switch, set "<monitoring.enabled>\" to \"false\".</monitoring.enabled> |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve Ethernet switch details

GET `/network/ethernet/switches/{name}`

**Introduced In:** 9.8

Retrieves the details of an Ethernet switch.

#### Related ONTAP commands

- `system switch ethernet show`

#### Learn more

- DOC /network/ethernet/switches

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Name |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| address | string | IP Address. |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| model | string | Model Number. |
| monitoring | monitoring | |
| name | string | Name. |
| network | string | Switch Network. |
| rcf_version | string | The switch reference configuration file (RCF) version. |

| Name | Type | Description |
| --- | --- | --- |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles.<br>Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "address": "string",
  "model": "string",
  "monitoring": {
    "reason": "string"
  },
  "name": "string",
  "network": "string",
  "rcf_version": "RCF NX9336C-FX2 v1.13 1-CLUSTER",
  "role": [
    "cluster+storage"
  ],
  "serial_number": "string",
  "snmp": {
    "user": "string",
    "version": "string"
  },
  "version": "string"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

monitoring

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable Health Monitoring. |
| monitored | boolean | Is Monitored. |
| reason | string | Reason For Not Monitoring. |

snmp

| Name | Type | Description |
|---|---|---|
| user | string | Community String or SNMPv3 Username. |
| version | string | SNMP Version. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an Ethernet switch

PATCH `/network/ethernet/switches/{name}`

**Introduced In:** 9.8

Configures monitoring settings for an Ethernet switch.

#### Related ONTAP commands

- `system switch ethernet modify`

#### Learn more

- DOC /network/ethernet/switches

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | path | True | Switch Name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| address | string | IP Address. |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| monitoring | monitoring | |

| Name | Type | Description |
|---|---|---|
| rcf_version | string | The switch reference configuration file (RCF) version. |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles.<br>Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

**Example request**

```
{
  "address": "string",
  "monitoring": {
    "reason": "string"
  },
  "rcf_version": "RCF NX9336C-FX2 v1.13 1-CLUSTER",
  "role": [
    "cluster+storage"
  ],
  "serial_number": "string",
  "snmp": {
    "user": "string",
    "version": "string"
  },
  "version": "string"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5636148 | Failed to discover SNMPv3 engine ID of the remote switch. |
| 5636149 | The "<snmp.user>\" must be a valid SNMPv3 user belonging to remote switch \"<name.>\" and must also be configured in ONTAP.</name.></snmp.user> |
| 5636152 | 0.0.0.0 cannot be specified as the IP address for the "<address.>\" parameter.</address.> |
| 8912900 | Invalid IP address. |
| 9044009 | SNMPv3 client feature is not available until the effective cluster version is ONTAP 9.3 or later. |
| 12517379 | Warning: Settings updated, but the IP "<address.>\" is not reachable. Verify that the address is valid, or check the network path.</address.> |
| 12517381 | Warning: Settings updated, but the SNMP validation request timed out. Verify that the "<snmp.user>\" parameter is valid.</snmp.user> |

| Error Code | Description |
| --- | --- |
| 12517383 | Warning: Settings updated, but the SNMPv3 validation request timed out. Verify that the "<snmp.user>\" parameter is valid (i.e., the SNMPv3 user exists in ONTAP and on the remote switch). If the \"<snmp.user>\" parameter is valid, verify that the SNMPv3 user’s credentials are the same both in ONTAP as well as in the remote switch. If a custom engine-id was provided for the SNMPv3 user, ensure it is the same as that of the remote switch.</snmp.user></snmp.user> |
| 12517411 | The input value for "<name.>\", \"<address.>\" and/or \"<network.>\" must match its discovered value.</network.></address.></name.> |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

monitoring

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable Health Monitoring. |
| monitored | boolean | Is Monitored. |
| reason | string | Reason For Not Monitoring. |

snmp

| Name | Type | Description |
|---|---|---|
| user | string | Community String or SNMPv3 Username. |
| version | string | SNMP Version. |

switch

Ethernet Switch REST API

| Name | Type | Description |
|---|---|---|
| address | string | IP Address. |
| discovered | boolean | Discovered By ONTAP CDP/LLDP<br>• readOnly: 1<br>• Introduced in: 9.8<br>• x-nullable: true |
| monitoring | monitoring | |
| rcf_version | string | The switch reference configuration file (RCF) version. |

| Name | Type | Description |
| --- | --- | --- |
| role | string | Switch role based on the interface configurations. Using '+' to combine multiple roles. Available role types are: cluster, multicluster, storage, multistorage, metrocluster. |
| serial_number | string | Serial Number. |
| snmp | snmp | |
| version | string | Software Version. |

_links

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Retrieve FC fabrics

## Retrieve FC fabrics

### Overview

The Fibre Channel (FC) fabric REST APIs provide read-only access to FC network information. This includes:

- the connections between the ONTAP cluster and the FC fabric,
- the switches that comprise the fabric, and
- the zones of the active zoneset of the fabric.

### Caching

Obtaining this information from the FC fabric can be time consuming. To allow the REST API to be more responsive, the APIs always return data from a cache that is updated asynchronously, but only on demand. Cache updates are triggered when the age of cached data exceeds the caller-specified maximum age as specified by the query parameter `cache.maximum_age`.

When a GET request initiates a cache refresh, the API attempts to wait for the update to complete before returning. If the cache cannot be updated before the return timeout (see query parameter `return_timeout`), the GET returns the currently cached data, but the cache update continues asynchronously. The caller may examine the returned property `cache.update_time` or `cache.age` to determine if the returned information is sufficiently fresh. If not, the caller should wait several seconds, then make a GET request again until the returned information is updated.

### Examples

Fibre Channel fabric data is typically large. The numbers of rows returned in the following examples has been edited to simplify reading.

#### Retrieving the Fibre Channel fabrics to which the cluster is connected

This example retrieves the names of the cluster's connected Fibre Channel fabrics. It also retrieves the cache timestamp properties so that the caller can verify the currency of the data.

```
# The API:
GET /api/network/fc/fabrics

# The call:
curl -X GET 'https://<mgmt-ip>/api/network/fc/fabrics?fields=cache' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "name": "10:00:aa:bb:cc:dd:ee:ff",
    "cache": {
      "update_time": "2022-02-07T21:21:29Z",
```

```
      "age": "PT1M16S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10%3A00%3Aaa%3Abb%3Acc%3Add%3Aee%3Aff"
      }
    }
  },
  {
    "name": "10:00:ff:ee:dd:cc:bb:aa",
    "cache": {
      "update_time": "2022-02-07T21:21:29Z",
      "age": "PT1M16S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10%3A00%3Aff%3Aee%3Add%3Acc%3Abb%3Aaa"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/fc/fabrics?fields=cache"
  }
}
}
```


---

#### Retrieving all switches of a Fibre Channel fabric

This example retrieves the switches of Fibre Channel fabric `10:00:aa:bb:cc:dd:ee:ff`.

```
# The API:
GET /api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches

# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches?fields=cache'
-H 'Accept: application/hal+json'
```

```
# The response:
{
"records": [
  {
    "wwn": "10:00:1a:1b:1c:1d:1e:1f",
    "cache": {
      "update_time": "2022-02-07T21:22:00Z",
      "age": "PT45S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:1a:1b:1c:1
d:1e:1f"
      }
    }
  },
  {
    "wwn": "10:00:2a:2b:2c:2d:2e:1f",
    "cache": {
      "update_time": "2022-02-07T21:22:00Z",
      "age": "PT45S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:2a:2b:2c:2
d:2e:1f"
      }
    }
  },
  {
    "wwn": "10:00:3a:3b:3c:3d:3e:3f",
    "cache": {
      "update_time": "2022-02-07T21:22:00Z",
      "age": "PT45S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:3a:3b:3c:3
d:3e:3f"
```

```
      }
    }
  },
  {
    "wwn": "10:00:4a:4b:4c:4d:4e:4f",
    "cache": {
      "update_time": "2022-02-07T21:22:00Z",
      "age": "PT45S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:4a:4b:4c:4
d:4e:4f"
      }
    }
  },
  {
    "wwn": "10:00:5a:5b:1a:5c:5d:5e",
    "cache": {
      "update_time": "2022-02-07T21:22:00Z",
      "age": "PT45S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:5a:5b:1a:5
c:5d:5e"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches?fields=cache"
  }
}
}
```

#### Retrieving all zones of the active zoneset of a Fibre Channel fabric

This example retrieves the zone of the active set of Fibre Channel fabric `10:00:aa:bb:cc:dd:ee:ff`.

```
# The API:
GET /api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones

# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones?fields=cache' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "name": "zone1",
    "cache": {
      "update_time": "2022-02-07T20:17:06Z",
      "age": "PT1H17M54S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones/zone1"
      }
    }
  },
  {
    "name": "zone2",
    "cache": {
      "update_time": "2022-02-07T20:17:06Z",
      "age": "PT1H17M54S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones/zone2"
      }
    }
  },
  {
    "name": "zone3",
    "cache": {
```

```
      "update_time": "2022-02-07T20:17:06Z",
      "age": "PT1H17M54S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones/zone3"
      }
    }
  },
  {
    "name": "zone4",
    "cache": {
      "update_time": "2022-02-07T20:17:06Z",
      "age": "PT1H17M54S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones/zone4"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/zones?fields=cache"
  }
}
}
```


#### Searching all Fibre Channel fabrics for a specific attached device identified by its WWPN

This example finds the Fibre Channel fabric, switch, and switch port to which the device with WWPN `50:0a:2a:2b:2c:2d:2e:2f` is attached. Note the use of the wildcard character in place of a fabric WWN in order to search all Fibre Channel fabrics.

```
# The API:
GET /api/network/fc/fabrics/*/switches
```

```
# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/fabrics/*/switches?fields=ports,cache&ports.attached_de
vice.wwpn=50:0a:2a:2b:2c:2d:2e:2f' -H 'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "fabric": {
      "name": "10:00:aa:bb:cc:dd:ee:ff",
      "_links": {
        "self": {
          "href":
"/api/network/fc/fabrics/10%3A00%3Aaa%3Abb%3Acc%3Add%3Aee%3Aff"
        }
      }
    },
    "wwn": "10:00:6a:6b:6c:6d:6e:6f",
    "ports": [
      {
        "slot": "0",
        "state": "online",
        "type": "f_port",
        "wwpn": "20:00:1a:1b:1c:1d:1e:1f",
        "attached_device": {
          "port_id": "0x999000",
          "wwpn": "50:0a:1a:1b:1c:1d:1e:1f"
        }
      },
      {
        "slot": "1",
        "state": "online",
        "type": "f_port",
        "wwpn": "20:01:2a:2b:1c:2d:2e:2f",
        "attached_device": {
          "port_id": "0x999100",
          "wwpn": "50:0a:2a:2b:2c:2d:2e:2f"
        }
      },
      {
        "slot": "2",
        "state": "offline",
        "type": "none",
        "wwpn": "20:02:3a:3b:3c:3d:3e:3f"
      },
```

```
      {
        "slot": "3",
        "state": "offline",
        "type": "f_port",
        "wwpn": "20:03:4a:4b:4c:4d:4e:4f",
        "attached_device": {
          "port_id": "0x999300",
          "wwpn": "50:0a:4a:4b:4c:4d:4e:4f"
        }
      },
      {
        "slot": "4",
        "state": "online",
        "type": "f_port",
        "wwpn": "20:04:5a:5b:5c:5d:5e:5f",
        "attached_device": {
          "port_id": "0x999400",
          "wwpn": "50:0a:5a:5b:5c:5d:5e:5f"
        }
      }
    ],
    "cache": {
      "update_time": "2022-02-07T21:57:29Z",
      "age": "PT4M49S",
      "is_current": "true"
    },
    "_links": {
      "self": {
        "href":
"/api/network/fc/fabrics/10:00:aa:bb:cc:dd:ee:ff/switches/10:00:6a:6b:6c:6
d:6e:6f"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/network/fc/fabrics/*/switches?fields=ports,cache&ports.attached_devi
ce.wwpn=50:0a:2a:2b:2c:2d:2e:2f"
  }
}
}
```

### Retrieve FC fabrics

GET `/network/fc/fabrics`

**Introduced In:** 9.11

Retrieves Fibre Channel fabrics.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connections`
- `zoneset`

#### Related ONTAP commands

- `network fcp topology show`
- `network fcp zone show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| zoneset.name | string | query | False | Filter by zoneset.name |
| connections.switch.wwn | string | query | False | Filter by connections.switch.wwn |
| connections.switch.port.wwpn | string | query | False | Filter by connections.switch.port.wwpn |
| connections.cluster_port.name | string | query | False | Filter by connections.cluster_port.name |
| connections.cluster_port.node.name | string | query | False | Filter by connections.cluster_port.node.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| connections.cluster_port.uuid | string | query | False | Filter by connections.cluster_port.uuid |
| connections.cluster_port.wwpn | string | query | False | Filter by connections.cluster_port.wwpn |
| name | string | query | False | Filter by name |
| cache.is_current | boolean | query | False | Filter by cache.is_current |
| cache.update_time | string | query | False | Filter by cache.update_time |
| cache.age | string | query | False | Filter by cache.age |
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fabric] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cache": {
        "age": "PT3M30S",
        "update_time": "string"
      },
      "connections": [
        {
          "cluster_port": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "0a",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
            "wwpn": "50:0a:11:22:33:44:55:66"
          },
          "switch": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "port": {
              "wwpn": "50:0a:a1:a2:a3:a4:a5:a6"
```

```
            },
            "wwn": "10:00:b1:b2:b3:b4:b4:b6"
          }
        }
      ],
      "name": "10:00:c1:c2:c3:c4:c5:c6",
      "zoneset": {
        "name": "zoneset1"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

node

The node on which the cluster Fibre Channel port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the cluster Fibre Channel port is located. |

cluster_port

The cluster Fibre Channel (FC) port that connects the FC fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the cluster Fibre Channel port. |
| node | node | The node on which the cluster Fibre Channel port is located. |
| uuid | string | The unique identifier of the cluster Fibre Channel port. |
| wwpn | string | The world wide port name (WWPN) of the cluster Fibre Channel port. |

port

The port of the Fibre Channel switch to which the cluster node port is connected.

| Name | Type | Description |
| --- | --- | --- |
| wwpn | string | The world wide port name (WWPN) of the Fibre Channel switch port. |

switch

The Fibre Channel switch to which the cluster node port is connected.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| port | port | The port of the Fibre Channel switch to which the cluster node port is connected. |
| wwn | string | The world-wide name (WWN) of the Fibre Channel switch to which the cluster node port is attached. |

connections

A connection between a cluster node Fibre Channel (FC) port and an FC switch port.

| Name | Type | Description |
| --- | --- | --- |
| cluster_port | cluster_port | The cluster Fibre Channel (FC) port that connects the FC fabric. |
| switch | switch | The Fibre Channel switch to which the cluster node port is connected. |

zoneset

The active Fibre Channel zoneset in the fabric.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the Fibre Channel zoneset. |

fabric

A Fibre Channel (FC) fabric REST object provides information about an FC network (fabric) connected to the cluster. Logically, the FC fabric also contains FC switches and the FC zones that comprise the active zoneset of the fabric. FC switch and zone information is not reported directly in the FC fabric REST object for reasons of scale and flexibility; they are found by querying the FC switches and FC zones REST endpoints.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cache | cache | Properties of Fibre Chanel fabric cache. |
| connections | array[connections] | An array of the connections between the cluster and the switches Fibre Channel fabric. |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |
| zoneset | zoneset | The active Fibre Channel zoneset in the fabric. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve FC fabric switches**

GET `/network/fc/fabrics/{fabric.name}/switches`

**Introduced In:** 9.11

Retrieves the Fibre Channel switches of a Fibre Channel fabric.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `ports`

#### Related ONTAP commands

- `network fcp topology show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fabric.name | string | path | True | The WWN of the primary switch of the Fibre Channel fabric. |
| name | string | query | False | Filter by name |
| release | string | query | False | Filter by release |
| vendor | string | query | False | Filter by vendor |
| wwn | string | query | False | Filter by wwn |
| cache.age | string | query | False | Filter by cache.age |
| cache.update_time | string | query | False | Filter by cache.update_time |
| cache.is_current | boolean | query | False | Filter by cache.is_current |
| domain_id | integer | query | False | Filter by domain_id<br>• Max value: 239<br>• Min value: 1 |
| ports.state | string | query | False | Filter by ports.state |
| ports.slot | string | query | False | Filter by ports.slot |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ports.type | string | query | False | Filter by ports.type |
| ports.attached_device.wwpn | string | query | False | Filter by ports.attached_device.wwpn |
| ports.attached_device.port_id | string | query | False | Filter by ports.attached_device.port_id |
| ports.wwpn | string | query | False | Filter by ports.wwpn |
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fc_switch] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cache": {
        "age": "PT3M30S",
        "update_time": "string"
      },
      "domain_id": 1,
      "fabric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "10:00:d1:d2:d3:d4:d5:d6"
      },
      "name": "switch1",
      "ports": [
        {
          "attached_device": {
            "port_id": 70400,
            "wwpn": "50:0a:21:22:23:24:25:26"
          },
          "slot": 1,
          "state": "online",
          "type": "string",
          "wwpn": "50:0a:31:32:33:34:35:36"
        }
      ],
      "release": "1.0.",
```

```
      "vendor": "vendor1",
      "wwn": "10:00:e1:e2:e3:e4:e5:e6"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5375053 | The Fibre Channel fabric specified by name in the request URI was not found in the FC fabric cache. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

fabric

A reference to a Fibre Channel fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |

attached_device

The Fibre Channel (FC) device attached to the FC switch port.

| Name | Type | Description |
| --- | --- | --- |
| port_id | string | The Fibre Channel port identifier of the attach device. |
| wwpn | string | The world-wide port name (WWPN) of the attached device. |

ports

A Fibre Channel switch port.

| Name | Type | Description |
| --- | --- | --- |
| attached_device | attached_device | The Fibre Channel (FC) device attached to the FC switch port. |
| slot | string | The slot of the Fibre Channel switch port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The state of the Fibre Channel switch port. |
| type | string | The type of the Fibre Channel switch port. |
| wwpn | string | The world wide port name (WWPN) of the Fibre Channel switch port. |

fc_switch

A Fibre Channel switch.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| cache | cache | Properties of Fibre Chanel fabric cache. |
| domain_id | integer | The domain identifier (ID) of the Fibre Channel (FC) switch. The domain ID is a unique identifier for the FC switch in the FC fabric. |
| fabric | fabric | A reference to a Fibre Channel fabric. |
| name | string | The logical name of the Fibre Channel switch. |
| ports | array[ports] | An array of the Fibre Channel (FC) switch's ports and their attached FC devices. |
| release | string | The firmware release of the Fibre Channel switch. |
| vendor | string | The vendor of the Fibre Channel switch. |
| wwn | string | The world-wide name (WWN) for the Fibre Channel switch. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC switch

GET `/network/fc/fabrics/{fabric.name}/switches/{wwn}`

**Introduced In:** 9.11

Retrieves a Fibre Channel switch.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `ports`

#### Related ONTAP commands

- `network fcp topology show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fabric.name | string | path | True | The WWN of the primary switch of the Fibre Channel fabric. |
| wwn | string | path | True | The WWN of the Fibre Channel switch. |
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cache | cache | Properties of Fibre Chanel fabric cache. |
| domain_id | integer | The domain identifier (ID) of the Fibre Channel (FC) switch. The domain ID is a unique identifier for the FC switch in the FC fabric. |
| fabric | fabric | A reference to a Fibre Channel fabric. |
| name | string | The logical name of the Fibre Channel switch. |

| Name | Type | Description |
| --- | --- | --- |
| ports | array[ports] | An array of the Fibre Channel (FC) switch's ports and their attached FC devices. |
| release | string | The firmware release of the Fibre Channel switch. |
| vendor | string | The vendor of the Fibre Channel switch. |
| wwn | string | The world-wide name (WWN) for the Fibre Channel switch. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cache": {
    "age": "PT3M30S",
    "update_time": "string"
  },
  "domain_id": 1,
  "fabric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "10:00:d1:d2:d3:d4:d5:d6"
  },
  "name": "switch1",
  "ports": [
    {
      "attached_device": {
        "port_id": 70400,
        "wwpn": "50:0a:21:22:23:24:25:26"
      },
      "slot": 1,
      "state": "online",
      "type": "string",
      "wwpn": "50:0a:31:32:33:34:35:36"
    }
  ],
  "release": "1.0.",
  "vendor": "vendor1",
  "wwn": "10:00:e1:e2:e3:e4:e5:e6"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5375053 | The Fibre Channel fabric specified by name in the request URI was not found in the FC fabric cache. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

fabric

A reference to a Fibre Channel fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |

attached_device

The Fibre Channel (FC) device attached to the FC switch port.

| Name | Type | Description |
| --- | --- | --- |
| port_id | string | The Fibre Channel port identifier of the attach device. |
| wwpn | string | The world-wide port name (WWPN) of the attached device. |

ports

A Fibre Channel switch port.

| Name | Type | Description |
| --- | --- | --- |
| attached_device | attached_device | The Fibre Channel (FC) device attached to the FC switch port. |
| slot | string | The slot of the Fibre Channel switch port. |

| Name | Type | Description |
|---|---|---|
| state | string | The state of the Fibre Channel switch port. |
| type | string | The type of the Fibre Channel switch port. |
| wwpn | string | The world wide port name (WWPN) of the Fibre Channel switch port. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve FC fabric zones

GET `/network/fc/fabrics/{fabric.name}/zones`

**Introduced In:** 9.11

Retrieves the zones of the active zoneset of a Fibre Channel fabric.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `members`

#### Related ONTAP commands

- `network fcp zone show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fabric.name | string | path | True | The WWN of the primary switch of the Fibre Channel fabric. |
| name | string | query | False | Filter by name |
| cache.age | string | query | False | Filter by cache.age |
| cache.is_current | boolean | query | False | Filter by cache.is_current |
| cache.update_time | string | query | False | Filter by cache.update_time |
| members.name | string | query | False | Filter by members.name |
| members.type | string | query | False | Filter by members.type |
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fc_zone] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cache": {
        "age": "PT3M30S",
        "update_time": "string"
      },
      "fabric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "10:00:d1:d2:d3:d4:d5:d6"
      },
      "members": [
        {
          "name": "10:00:12:34:56:78:9a:bc",
          "type": "port_name"
        }
      ],
      "name": "zone1"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5375053 | The Fibre Channel fabric specified by name in the request URI was not found in the FC fabric cache. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

fabric

A reference to a Fibre Channel fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |

members

A Fibre Channel zone member.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The identifying property value of the zone member. The format of this value depends on the member type:<br><br>• `port_id`: A zero-filled 6-digit hexadecimal value with a 0x prefer. Example: 0x0000A0.<br>• `port_name`: A world-wide name. Example: 10:00:12:34:56:78:9a:bc.<br>• `domain_id_port`: A domain ID and a port ID as decimal integers separated by a slash. Example: 1/2.<br>• `node_name`: A world-wide name. Example: 10:00:11:22:33:44:55:66.<br>• `fabric_port_name`: A world-wide name. Example: 10:00:ab:cd:ef:12:34:56.<br><br>#<br><br>The following types might not report a name:<br><br>• `interface`<br>• `domain_interface`<br>• `ip_address`<br>• `symbolic_node_name`<br>• `device_alias` |
| type | string | The type of Fibre Channel zone member. This value should be used to interpret the contents of the `name` property. |

fc_zone

A Fibre Channel zone.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| cache | cache | Properties of Fibre Chanel fabric cache. |
| fabric | fabric | A reference to a Fibre Channel fabric. |
| members | array[members] | An array of Fibre Channel zone members. |
| name | string | The name of the Fibre Channel zone. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC fabric zone

GET `/network/fc/fabrics/{fabric.name}/zones/{name}`

**Introduced In:** 9.11

Retrieves a zone of the active zoneset of a Fibre Channel fabric.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `members`

#### Related ONTAP commands

- `network fcp zone show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fabric.name | string | path | True | The WWN of the primary switch of the Fibre Channel fabric. |
| name | string | path | True | The name of a zone in the active zoneset the Fibre Channel fabric. |
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cache | cache | Properties of Fibre Chanel fabric cache. |

| Name | Type | Description |
|---|---|---|
| fabric | fabric | A reference to a Fibre Channel fabric. |
| members | array[members] | An array of Fibre Channel zone members. |
| name | string | The name of the Fibre Channel zone. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cache": {
    "age": "PT3M30S",
    "update_time": "string"
  },
  "fabric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "10:00:d1:d2:d3:d4:d5:d6"
  },
  "members": [
    {
      "name": "10:00:12:34:56:78:9a:bc",
      "type": "port_name"
    }
  ],
  "name": "zone1"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5375053 | The Fibre Channel fabric specified by name in the request URI was not found in the FC fabric cache. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

fabric

A reference to a Fibre Channel fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |

members

A Fibre Channel zone member.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The identifying property value of the zone member. The format of this value depends on the member type:<br><br>• `port_id`: A zero-filled 6-digit hexadecimal value with a 0x prefer. Example: 0x0000A0.<br>• `port_name`: A world-wide name. Example: 10:00:12:34:56:78:9a:bc.<br>• `domain_id_port`: A domain ID and a port ID as decimal integers separated by a slash. Example: 1/2.<br>• `node_name`: A world-wide name. Example: 10:00:11:22:33:44:55:66.<br>• `fabric_port_name`: A world-wide name. Example: 10:00:ab:cd:ef:12:34:56.<br><br>#<br><br>The following types might not report a name:<br><br>• `interface`<br>• `domain_interface`<br>• `ip_address`<br>• `symbolic_node_name`<br>• `device_alias` |
| type | string | The type of Fibre Channel zone member. This value should be used to interpret the contents of the `name` property. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC fabric

GET `/network/fc/fabrics/{name}`

**Introduced In:** 9.11

Retrieves a Fibre Channel fabric.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connections`
- `zoneset`

#### Related ONTAP commands

- `network fcp topology show`
- `network fcp zone show`

#### Learn more

- DOC /network/fc/fabrics

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | The world wide name (WWN) of the primary switch of the Fibre Channel fabric. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cache.maximum_age | string | query | False | The maximum age of data in the Fibre Channel fabric cache before it should be refreshed from the fabric. The default is 15 minutes.<br>• Default value: 1<br>• format: iso8601 |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cache | cache | Properties of Fibre Chanel fabric cache. |
| connections | array[connections] | An array of the connections between the cluster and the switches Fibre Channel fabric. |
| name | string | The world wide name (WWN) of the primary switch of the Fibre Channel (FC) fabric. This is used as a unique identifier for the FC fabric. |
| zoneset | zoneset | The active Fibre Channel zoneset in the fabric. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cache": {
    "age": "PT3M30S",
    "update_time": "string"
  },
  "connections": [
    {
      "cluster_port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
        "wwpn": "50:0a:11:22:33:44:55:66"
      },
      "switch": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "port": {
          "wwpn": "50:0a:a1:a2:a3:a4:a5:a6"
        },
        "wwn": "10:00:b1:b2:b3:b4:b4:b6"
      }
    }
  ],
  "name": "10:00:c1:c2:c3:c4:c5:c6",
  "zoneset": {
    "name": "zoneset1"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cache

Properties of Fibre Chanel fabric cache.

| Name | Type | Description |
|---|---|---|
| age | string | The age of the Fibre Channel fabric data cache retrieved. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. The value is in ISO 8601 duration format. |
| is_current | boolean | A boolean that indicates if the retrieved data is current relative to the `cache.maximum_age` value of the request. A value of `true` indicates that the data is no older than the requested maximum age. A value of `false` indicates that the data is older than the requested maximum age; if more current data is required, the caller should wait for some time for the cache update to complete and query the data again. |

| Name | Type | Description |
| --- | --- | --- |
| update_time | string | The date and time at which the Fibre Channel fabric data cache retrieved was last updated. If the FC fabric data cache has not been fully updated for a newly discovered fabric, or a fabric that has been re-discovered after being purged, a value for this property will not be retrieved. |

node

The node on which the cluster Fibre Channel port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the cluster Fibre Channel port is located. |

cluster_port

The cluster Fibre Channel (FC) port that connects the FC fabric.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the cluster Fibre Channel port. |
| node | node | The node on which the cluster Fibre Channel port is located. |
| uuid | string | The unique identifier of the cluster Fibre Channel port. |
| wwpn | string | The world wide port name (WWPN) of the cluster Fibre Channel port. |

port

The port of the Fibre Channel switch to which the cluster node port is connected.

| Name | Type | Description |
| --- | --- | --- |
| wwpn | string | The world wide port name (WWPN) of the Fibre Channel switch port. |

switch

The Fibre Channel switch to which the cluster node port is connected.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| port | port | The port of the Fibre Channel switch to which the cluster node port is connected. |
| wwn | string | The world-wide name (WWN) of the Fibre Channel switch to which the cluster node port is attached. |

connections

A connection between a cluster node Fibre Channel (FC) port and an FC switch port.

| Name | Type | Description |
| --- | --- | --- |
| cluster_port | cluster_port | The cluster Fibre Channel (FC) port that connects the FC fabric. |
| switch | switch | The Fibre Channel switch to which the cluster node port is connected. |

zoneset

The active Fibre Channel zoneset in the fabric.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the Fibre Channel zoneset. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FC network interfaces

### Manage FC network interfaces

#### Overview

Fibre Channel (FC) interfaces are the logical endpoints for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol (FCP) or NVMe over FC (NVMe/FC).

The FC interface REST API allows you to create, delete, update, and discover FC interfaces, and obtain status information for FC interfaces.

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH request that relocates an interface. You can identify the port by supplying either the node and port names or the port UUID.

#### Performance monitoring

Performance of an FC interface can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the performance of an FC interface in terms of IOPS, latency, and throughput. The `metric.*` properties denote an average, whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Interface placement recommendations

The FC interface REST API can also recommend the placement (cluster nodes and FC ports) for FC interfaces for a new or existing SVM as well as evaluate caller-proposed locations for FC interfaces. This functionality is available to cluster administrators only and is accessed using `GET /network/fc/interfaces` with the `recommend` family of query parameters.

The query parameter `recommend.data_protocol` is required when getting recommendations or evaluating caller-proposed locations for FC interfaces. It identifies the type of FC interfaces to recommend. Other

`recommend` query parameters are optional and are used to modify the recommendation algorithm.

If an SVM is supplied using the query parameter `recommend.svm.name` and/or `recommend.svm.uuid`, existing FC interfaces are considered as part of the overall solution and only additionally recommended interfaces are returned. If no SVM is supplied, recommendations are returned for a new SVM.

FC fabrics connected to the cluster are discovered by the API. By default, FC interfaces are placed and evaluated for each fabric. The query parameter `recommend.fabrics.name` can be used to identify specific FC fabrics to use.

Cluster nodes supporting FC fabric connections for the specific data protocol are discovered by the API. By default, FC interfaces are placed all supported cluster nodes. Either query parameter `recommend.nodes.name` or `recommend.nodes.uuid` can be used to identify specific cluster nodes to use.

FC interfaces for the FC-NVMe data protocol are limited to two (2) interfaces per cluster node with a maximum of four (4) nodes, within a single SVM.

Placement recommendations are best effort and limited by the information available. In situations where an optimum configuration cannot be produced, the API returns the recommendations it can along with messages describing how the caller might improve the configuration. These messages are produced by evaluating the calculated FC interface layout against best practices.

The same best practice evaluation can be applied to a caller-proposed configuration by using the query parameter `recommend.proposed.locations.port.uuid` to specify the locations for proposed FC interfaces. When this query parameter is supplied, the best practice evaluation is performed using the proposed interface locations and messages are produced describing how the caller might improve the configuration.

#### Examples

### Creating an FC interface using the port node and name to identify the location

This example uses the `return_records` query parameter to retrieve the newly created FC interface in the POST response.

```
# The API:
POST /api/network/fc/interfaces

# The call:
curl -X POST 'https://<mgmt-
ip>/api/network/fc/interfaces?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "lif1",
"location": { "home_port": { "name": "0a", "home_node": { "name": "node1"
} } }, "data_protocol": "fcp" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
```

```
      "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"
        }
      }
    },
    "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",
    "name": "lif1",
    "location": {
      "home_node": {
        "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "home_port": {
        "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "node": {
        "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "port": {
        "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
```

```
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-
005056bba0e0"
          }
        }
      }
    },
    "enabled": true,
    "state": "down",
    "data_protocol": "fcp",
    "wwpn": "20:04:00:50:56:bb:a0:e0",
    "wwnn": "20:00:00:50:56:bb:a0:e0",
    "port_address": "9da2cb1",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-
005056bba0e0"
      }
    }
  }
]
}
```


#### Creating an FC interface using the port UUID to identify the location

This example uses the `return_records` query parameter to retrieve the newly created FC interface in the POST response.

```
# The API:
POST /api/network/fc/interfaces

# The call:
curl -X POST 'https://<mgmt-
ip>/api/network/fc/interfaces?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm3" }, "name": "lif2",
"location": { "home_port": { "uuid": "24bb636a-db83-11e8-9a49-
005056bb1ec6" } }, "data_protocol": "fc_nvme" }'

# The response:
```

```
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
      "name": "svm3",
      "_links": {
        "self": {
          "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
        }
      }
    },
    "uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
    "name": "lif2",
    "location": {
      "home_node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
          }
        }
      },
      "home_port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {
          "name": "node3"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
          }
        }
      },
      "node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
```

```
          }
        }
      },
      "port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {
          "name": "node3"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
          }
        }
      }
    },
    "enabled": true,
    "state": "down",
    "data_protocol": "fc_nvme",
    "wwpn": "20:05:00:50:56:bb:a0:e0",
    "wwnn": "20:02:00:50:56:bb:a0:e0",
    "port_address": "612e202b",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-
005056bba0e0"
      }
    }
  }
]
}
```


---

#### Retrieving all properties for all FC interfaces

This example uses the `fields` query parameter to retrieve all properties.

```
# The API:
GET /api/network/fc/interfaces

# The call:
curl -X GET 'https://<mgmt-ip>/api/network/fc/interfaces?fields=*' -H
'Accept: application/hal+json'
```

---

```
# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
      "name": "svm3",
      "_links": {
        "self": {
          "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
        }
      }
    },
    "uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
    "name": "lif2",
    "location": {
      "home_node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
          }
        }
      },
      "home_port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {
          "name": "node3"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
          }
        }
      },
      "node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
```

```
          }
        }
      },
      "port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {
          "name": "node3"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
          }
        }
      }
    },
    "enabled": true,
    "state": "down",
    "data_protocol": "fc_nvme",
    "wwpn": "20:05:00:50:56:bb:a0:e0",
    "wwnn": "20:02:00:50:56:bb:a0:e0",
    "port_address": "612e202b",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-
005056bba0e0"
      }
    }
  },
  {
    "svm": {
      "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"
        }
      }
    },
    "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",
    "name": "lif1",
    "location": {
      "home_node": {
        "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
        "name": "node1",
```

```
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "home_port": {
        "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "node": {
        "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-
005056bba0e0"
          }
        }
      },
      "port": {
        "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-
005056bba0e0"
          }
        }
      }
    },
    "enabled": true,
    "state": "down",
```

```
    "data_protocol": "fcp",
    "wwpn": "20:04:00:50:56:bb:a0:e0",
    "wwnn": "20:00:00:50:56:bb:a0:e0",
    "port_address": "9da2cb1",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-
005056bba0e0"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/fc/interfaces?fields=*"
  }
}
}
```


---

### Retrieving a list of selected FC interfaces

This example uses property query parameters to retrieve FC interfaces configured for the FC Protocol that are set to *up*.

```
# The API:
GET /api/network/fc/interfaces

# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/interfaces?data_protocol=fcp&state=up' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"
        }
      }
    },
    "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",
    "name": "lif1",
    "state": "up",
    "data_protocol": "fcp",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-
005056bba0e0"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/fc/interfaces?data_protocol=fcp&state=up"
  }
}
}
```

#### Retrieving a specific FC interface

```
# The API:
GET /api/network/fc/interfaces/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/network/fc/interfaces/cdeb5591-dec9-
11e8-a733-005056bba0e0' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
  "name": "svm3",
  "_links": {
    "self": {
      "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
    }
  }
},
"uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
"name": "lif2",
"location": {
  "home_node": {
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-005056bb1ec6"
      }
    }
  },
  "home_port": {
    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
      }
    }
  },
  "node": {
```

```
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-005056bb1ec6"
      }
    }
  },
  "port": {
    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
      }
    }
  }
},
"enabled": true,
"state": "down",
"data_protocol": "fc_nvme",
"wwpn": "20:05:00:50:56:bb:a0:e0",
"wwnn": "20:02:00:50:56:bb:a0:e0",
"port_address": "612e202b",
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
```

```
    "write": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-
005056bba0e0"
  }
}
}
```


#### Disabling an FC interface

When updating certain properties of an FC interface, the interface must first be disabled using the following:

```
# The API:
PATCH /api/network/fc/interfaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-
11e8-a733-005056bba0e0' -H 'Accept: application/hal+json' -d '{ "enabled":
false }'
```

---

#### Moving an FC interface to a new node and port

To move an FC interface to another node or port, the destination FC port must be specified in a PATCH request. Either the port UUID or node and port names can be used to identify the port.

Note that only FC interfaces configured for the FC Protocol can be moved. FC interfaces configured for NVMe/FC cannot be moved. The interface must also be set to the disabled state before being moved.

```
# The API:
PATCH /api/network/fc/interfaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-
11e8-a733-005056bba0e0' -H 'Accept: application/hal+json' -d '{
"location": { "home_port": { "uuid": "a1dc7aa5-db83-11e8-9ef7-
005056bbbbcc" } } }'
```

---

#### Deleting an FC interface

```
# The API:
DELETE /api/network/fc/interfaces/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-
11e8-a733-005056bba0e0' -H 'Accept: application/hal+json'
```

---

#### Recommending interface locations for a new SVM

This example gets recommendations for FCP network interfaces for a new SVM.

```
# The API
```

```
GET /api/network/fc/interfaces

# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/interfaces?recommend.data_protocol=fcp&fields=*' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "location": {
      "home_port": {
        "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-
005056bba0e0"
          }
        }
      }
    },
    "data_protocol": "fcp",
    "comment": "fabric: 55:0e:b1:a0:20:40:80:00"
  },
  {
    "location": {
      "home_port": {
        "uuid": "ad7d3915-db82-11e8-b36d-005056bb982e",
        "name": "0a",
        "node": {
          "name": "node2"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/ad7d3915-db82-11e8-b36d-
005056bb982e"
          }
        }
      }
    },
    "data_protocol": "fcp",
```

```
    "comment": "fabric: 55:0e:b1:a0:20:40:80:00"
  },
  {
    "location": {
      "home_port": {
        "uuid": "300c1dfd-db82-11e8-bd46-005056bba0e0",
        "name": "0b",
        "node": {
          "name": "node1"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/300c1dfd-db82-11e8-bd46-
005056bba0e0"
          }
        }
      }
    },
    "data_protocol": "fcp",
    "comment": "fabric: 55:0e:b1:a0:20:40:80:01"
  }
],
"num_records": 3,
"recommend": {
  "messages": [
    {
      "message": "The following network ports are disabled (cluster
node:port): node2:0b",
      "code": "5375959",
      "arguments": [
        {
          "message": "node2:0b"
        }
      ],
      "severity": "informational"
    }
  ]
},
"_links": {
  "self": {
    "href":
"/api/network/fc/interfaces?recommend.data_protocol=fcp&fields=*"
  }
}
}
```

#### Proposing interface locations for a new SVM

This example requests that caller-proposed locations for FC-NVMe interfaces on two nodes be evaluated.

```
# The API
GET /api/network/fc/interfaces

# The call:
curl -X GET 'https://<mgmt-
ip>/api/network/fc/interfaces?&recommend.data_protocol=fc_nvme&recommend.p
roposed.locations.port.uuid=300c2786-db82-11e8-bd46-005056bba0e0,ad7d47d6-
db82-11e8-b36d-005056bb982e&fields=*' -H 'Accept: application/hal+json'

# The response:
{
"records": [
],
"num_records": 0,
"recommend": {
  "messages": [
    {
      "message": "Cluster node \"node1\" does not have the requested
number (1) of network interfaces for the following Fibre Channel fabrics:
55:0e:b1:a0:20:40:80:07.",
      "code": "5375969",
      "arguments": [
        {
          "message": "node1"
        },
        {
          "message": "1"
        },
        {
          "message": "55:0e:b1:a0:20:40:80:07"
        }
      ],
      "severity": "warning"
    },
    {
      "message": "Cluster node \"node2\" does not have the requested
number (1) of network interfaces for the following Fibre Channel fabrics:
55:0e:b1:a0:20:40:80:07.",
      "code": "5375969",
      "arguments": [
        {
          "message": "node2"
```

```
        },
        {
          "message": "1"
        },
        {
          "message": "55:0e:b1:a0:20:40:80:07"
        }
      ],
      "severity": "warning"
    },
    {
      "message": "The SVM cannot be reached from all of the Fibre Channel
(FC) fabrics to which the cluster is connected. Cluster FC fabrics:
55:0e:b1:a0:20:40:80:06, 55:0e:b1:a0:20:40:80:07. SVM FC fabrics:
55:0e:b1:a0:20:40:80:06.",
      "code": "5375970",
      "arguments": [
        {
          "message": "55:0e:b1:a0:20:40:80:06, 55:0e:b1:a0:20:40:80:07"
        },
        {
          "message": "55:0e:b1:a0:20:40:80:06"
        }
      ],
      "severity": "warning"
    },
    {
      "message": "The SVM is configured for 1 Fibre Channel (FC) fabrics.
The preferred configuration is 2 FC fabrics.",
      "code": "5375973",
      "arguments": [
        {
          "message": "1"
        },
        {
          "message": "2"
        }
      ],
      "severity": "warning"
    }
  ]
},
"_links": {
  "self": {
    "href":
"/api/network/fc/interfaces?recommend.data_protocol=fc_nvme&recommend.node
```

```
s.name=node1,node2&recommend.proposed.locations.port.uuid=300c2786-db82-
11e8-bd46-005056bba0e0,ad7d47d6-db82-11e8-b36d-005056bb982e&fields=*"
  }
}
}
```

---

### Retrieve FC interfaces

GET `/network/fc/interfaces`

**Introduced In:** 9.6

Retrieves FC interfaces.

#### Related ONTAP commands

- `network interface show`
- `vserver fcp interface show`

#### Learn more

- DOC /network/fc/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| recommend.svm.name | string | query | False | The name of an existing SVM for which to recommend FC interfaces. If neither query parameter `recommend.svm.name` nor `recommend.svm.uuid` are provided when recommending FC interfaces, recommendations are made for a new, non-existent SVM.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.svm.uuid | string | query | False | The UUID of an existing SVM for which to recommend FC interfaces. If neither query parameter `recommend.svm.name` nor `recommend.svm.uuid` are provided when recommending FC interfaces, recommendations are made for a new, non-existent SVM.<br>• Introduced in: 9.11 |
| recommend.data_protocol | string | query | False | The target data protocol for which to recommend FC interfaces. This query parameter is required when recommending FC interfaces.<br>• Introduced in: 9.11<br>• enum: ["fcp", "fc_nvme"] |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.interfaces_per_node | integer | query | False | The target number of interfaces per FC fabric per cluster node when recommending FC interfaces. If this query parameter is not provided, the default value is one (1).<br>• Introduced in: 9.11<br>• Default value: 1<br>• Max value: 4<br>• Min value: 1 |
| recommend.fabrics.name | array[string] | query | False | The names of the FC fabrics to target when recommending FC interfaces. If this query parameter is not provided, FC interfaces will be recommended for each FC fabric discovered.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.nodes.name | array[string] | query | False | The names of the cluster nodes to target when recommending FC interfaces. If this query parameter is not provided, FC interfaces will be recommended for all nodes that have FC network ports supporting target data protocol. Query parameters `recommend.nodes.name` and `recommend.nodes.uuid` are mutually exclusive; use either names or UUIDs to identify cluster nodes.<br>• Introduced in: 9.11 |
| recommend.nodes.uuid | array[string] | query | False | The UUIDs of the cluster nodes to target when recommending FC interfaces. If this query parameter is not provided, FC interfaces will be recommended for all nodes that have FC network ports supporting target data protocol. Query parameters `recommend.nodes.name` and `recommend.nodes.uuid` are mutually exclusive; use either names or UUIDs to identify cluster nodes.<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.proposed.locations.port.uuid | array[string] | query | False | The UUIDs of the FC ports on which FC interfaces are proposed. A UUID may be supplied multiple times to proposed multiple FC interfaces. FC ports must be enabled, support the target data protocol and be reporting an FC fabric.<br><br>• Introduced in: 9.11 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| state | string | query | False | Filter by state |
| name | string | query | False | Filter by name |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br>• Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.8 |
| location.port.name | string | query | False | Filter by location.port.name |
| location.port.node.name | string | query | False | Filter by location.port.node.name |
| location.port.uuid | string | query | False | Filter by location.port.uuid |
| location.home_port.name | string | query | False | Filter by location.home_port.name<br>• Introduced in: 9.8 |
| location.home_port.node.name | string | query | False | Filter by location.home_port.node.name<br>• Introduced in: 9.8 |
| location.home_port.uuid | string | query | False | Filter by location.home_port.uuid<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| location.home_node.name | string | query | False | Filter by location.home_node.name<br>• Introduced in: 9.8 |
| location.home_node.uuid | string | query | False | Filter by location.home_node.uuid<br>• Introduced in: 9.8 |
| location.is_home | boolean | query | False | Filter by location.is_home<br>• Introduced in: 9.8 |
| location.node.name | string | query | False | Filter by location.node.name |
| location.node.uuid | string | query | False | Filter by location.node.uuid |
| wwpn | string | query | False | Filter by wwpn |
| enabled | boolean | query | False | Filter by enabled |
| port_address | string | query | False | Filter by port_address |
| comment | string | query | False | Filter by comment |
| data_protocol | string | query | False | Filter by data_protocol |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.8 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.8 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.8 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.8 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.8 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.8 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br>• Introduced in: 9.8 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.8 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.8 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |
| wwnn | string | query | False | Filter by wwnn |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| recommend | recommend | Response properties specific to the FC interface placement functionality. See the *Interface placement recommendations* section of DOC /network/fc/interfaces |

| Name | Type | Description |
| --- | --- | --- |
| records | array[fc_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "recommend": {
    "messages": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": 5375959,
        "message": "Network ports are disabled.",
        "severity": "informational"
      }
    ]
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "data_protocol": "string",
      "location": {
        "home_node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
```

```
      "home_port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "0a",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "metric": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
```

```
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "fc_lif1",
      "port_address": "5060F",
      "state": "string",
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "bce9827d-4d8f-60af-c771-6e8e9af2c6f0",
      "wwnn": "20:00:00:50:56:b4:13:01",
```

```
        "wwpn": "20:00:00:50:56:b4:13:a8"
      }
    ]
  }
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5374938 | Query parameter `recommend.data_protocol` is required when any other `recommend` query parameters are specified. |
| 5374939 | Query parameter `recommend.fabrics.name` specifies a duplicate FC fabric. |
| 5374940 | Query parameter `recommend.nodes.name` or `recommend.nodes.uuid` specifies a duplicate node. |
| 5375953 | Query parameter `recommend.fabrics.name` specifies an FC fabric to which no FC port is connected. |
| 5375954 | Query parameter `recommend.nodes.name` or `recommend.nodes.uuid` specifies a node that does not support the specified data protocol. |
| 5375955 | Query parameter `recommend.proposed.locations.port.uuid` does not specify a valid port or specifies a port that is support the specified data protocol. |
| 5375956 | Query parameter `recommend.proposed.locations.port.uuid` specifies a port that is not located on a node specified by query parameter `recommend.nodes.name` or `recommend.nodes.uuid`. If the nodes are to be constrained, only ports on those nodes may be proposed. |
| 5375957 | Query parameter `recommend.proposed.locations.port.uuid` specifies a port that is disabled. |

| Error Code | Description |
| --- | --- |
| 5375958 | Query parameter `recommend.proposed.locations.port.uuid` specifies a port that is not reporting a connected FC fabric. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

fc_interface_recommend_message

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | The message substitution arguments. |

| Name | Type | Description |
| --- | --- | --- |
| code | string | The message code. Possible messages:<br><br>ONTAP Error Response Codes<br><br>Error Code — Description<br><br>5375959 — Network ports are disabled.<br><br>5375960 — Network ports are enabled, but not reporting a connected FC fabric.<br><br>5375961 — The limit for the number of FC network interfaces on a cluster node has been reached.<br><br>5375962 — The limit for the number of FC network interfaces on a port has been reached.<br><br>5375963 — An HA pair of cluster nodes has a discrepancy in the presence of FC ports.<br><br>5375964 — An HA pair of cluster nodes has a discrepancy in support for an FC data protocol.<br><br>5375965 — An HA pair of cluster nodes cannot be reached from the same FC fabrics.<br><br>5375966 — A cluster node cannot be reached from all of the FC fabrics from which other cluster nodes with FC interfaces in the SVM can be reached.<br><br>5375967 — The limit for the number of FC network interfaces on a cluster node has been exceeded.<br><br>5375968 — The limit for the number of FC network interfaces on an FC port has been exceeded.<br><br>5375969 — The requested number of network interfaces per FC fabric per cluster node has not been achieved. |
| | | 5375970 — The SVM cannot be |

| Name | Type | Description |
| --- | --- | --- |
| message | string | The message text. |
| severity | string | The severity of the message. Message severities are as follows:<br>• `error` - Messages reporting problems that must be corrected before creating the FC network interfaces.<br>• `warning` - Messages indicating issues that need rectifying in order to achieve an optimal configuration.<br>• `informational` - Messages providing relevant information for consideration. |

recommend

Response properties specific to the FC interface placement functionality. See the *Interface placement recommendations* section of DOC /network/fc/interfaces

| Name | Type | Description |
| --- | --- | --- |
| messages | array[fc_interface_recommend_message] | Messages describing the results of a FC network interface placement operation or evaluation of caller-proposed locations. |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

home_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

node

The node on which the FC port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the FC port is located. |

home_port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID,

or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.

`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship.

| Name | Type | Description |
|---|---|---|
| home_node | home_node | |
| home_port | home_port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| is_home | boolean | Indicates if the FC interface is currently on its home node. |
| node | node | |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fc_interface

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the cluster node and port names or the port UUID.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| data_protocol | string | The data protocol for which the FC interface is configured. Required in POST. |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |

| Name | Type | Description |
| --- | --- | --- |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an FC interface**

POST `/network/fc/interfaces`

**Introduced In:** 9.6

Creates an FC interface.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the FC interface.
- `name` - Name of the FC interface.
- `location.port.uuid` or both `location.port.name` and `location.port.node.name` - FC port on which to create the FC interface.
- `data_protocol` - Data protocol for the FC interface.

#### Default property values

If not specified in POST, the following default property values are assigned.

- `enabled` - *true*

#### Related ONTAP commands

- `network interface create`

#### Learn more

- DOC /network/fc/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| data_protocol | string | The data protocol for which the FC interface is configured. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br><br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

### Example request

```
{
  "comment": "string",
  "data_protocol": "string",
  "location": {
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "fc_lif1",
  "port_address": "5060F",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "bce9827d-4d8f-60af-c771-6e8e9af2c6f0",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| recommend | recommend | Response properties specific to the FC interface placement functionality. See the *Interface placement recommendations* section of DOC /network/fc/interfaces |
| records | array[fc_interface] | |

**Example response**


```
{
  "num_records": 1,
  "recommend": {
    "messages": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": 5375959,
        "message": "Network ports are disabled.",
        "severity": "informational"
      }
    ]
  },
  "records": [
    {
      "comment": "string",
      "data_protocol": "string",
      "location": {
        "home_node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "home_port": {
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "port": {
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
```

```
      },
      "name": "fc_lif1",
      "port_address": "5060F",
      "state": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "bce9827d-4d8f-60af-c771-6e8e9af2c6f0",
      "wwnn": "20:00:00:50:56:b4:13:01",
      "wwpn": "20:00:00:50:56:b4:13:a8"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966140 | An interface with the same name already exists. |
| 1966217 | The specified port is not valid on the node provided. |
| 2621462 | The supplied SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373966 | A Fibre Channel interface with the *fcp* protocol cannot be created in an SVM that is configured for NVMe. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |

| Error Code | Description |
| --- | --- |
| 5374102 | The specified Fibre Channel interface cannot be created because the Fibre Channel adapter is down. Bring the adapter up and try again. |
| 5374871 | The Fibre Channel port identified by the specified UUID does not refer to the same port as that identified by the specified node name and/or port name. |
| 5374872 | If either `location.port.node.name` or `location.port.name` is supplied, both properties must be supplied. |
| 5374873 | The Fibre Channel port must be specified using either `location.port.uuid` or `location.port.node.name` and `location.port.name`. |
| 5375057 | An FC port with the provided UUID does not exist. |
| 72089652 | An NVMe service must be created before creating a Fibre Channel interface using the NVMe over FC data protocol. |
| 72089672 | The specified Fibre Channel port does not support the NVMe over FC data protocol. |
| 72089900 | A Fibre Channel interface with the *fc_nvme* protocol cannot be created in an SVM that is configured for a SAN protocol. |
| 72089901 | NVMe over Fabrics does not support more than 2 LIFs per node per SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

home_node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

home_port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC

network.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.

`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship.

| Name | Type | Description |
|---|---|---|
| home_node | home_node | |
| home_port | home_port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| is_home | boolean | Indicates if the FC interface is currently on its home node. |
| node | node | |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fc_interface

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the cluster node and port names or the port UUID.

| Name | Type | Description |
|---|---|---|
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| data_protocol | string | The data protocol for which the FC interface is configured. Required in POST. |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br><br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br><br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

fc_interface_recommend_message

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | The message substitution arguments. |

| Name | Type | Description |
| --- | --- | --- |
| code | string | The message code. Possible messages:<br><br>ONTAP Error Response Codes<br><br>Error Code — Description<br><br>5375959 — Network ports are disabled.<br><br>5375960 — Network ports are enabled, but not reporting a connected FC fabric.<br><br>5375961 — The limit for the number of FC network interfaces on a cluster node has been reached.<br><br>5375962 — The limit for the number of FC network interfaces on a port has been reached.<br><br>5375963 — An HA pair of cluster nodes has a discrepancy in the presence of FC ports.<br><br>5375964 — An HA pair of cluster nodes has a discrepancy in support for an FC data protocol.<br><br>5375965 — An HA pair of cluster nodes cannot be reached from the same FC fabrics.<br><br>5375966 — A cluster node cannot be reached from all of the FC fabrics from which other cluster nodes with FC interfaces in the SVM can be reached.<br><br>5375967 — The limit for the number of FC network interfaces on a cluster node has been exceeded.<br><br>5375968 — The limit for the number of FC network interfaces on an FC port has been exceeded.<br><br>5375969 — The requested number of network interfaces per FC fabric per cluster node has not been achieved. |
| | | 5375970 — The SVM cannot be |

| Name | Type | Description |
| --- | --- | --- |
| message | string | The message text. |
| severity | string | The severity of the message. Message severities are as follows:<br>• `error` - Messages reporting problems that must be corrected before creating the FC network interfaces.<br>• `warning` - Messages indicating issues that need rectifying in order to achieve an optimal configuration.<br>• `informational` - Messages providing relevant information for consideration. |

recommend

Response properties specific to the FC interface placement functionality. See the *Interface placement recommendations* section of DOC /network/fc/interfaces

| Name | Type | Description |
| --- | --- | --- |
| messages | array[fc_interface_recommend_message] | Messages describing the results of a FC network interface placement operation or evaluation of caller-proposed locations. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve historical performance metrics for an FC interface

GET `/network/fc/interfaces/{fc_interface.uuid}/metrics`

**Introduced In:** 9.14

Retrieves historical performance metrics for a Fibre Channel interface.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| status | string | query | False | Filter by status |
| uuid | string | query | False | Filter by uuid |
| duration | string | query | False | Filter by duration |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| fc_interface.uuid | string | path | True | The unique identifier of the Fibre Channel interface. |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the FC interface. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve historical performance metrics for an FC interface for a specific time

GET `/network/fc/interfaces/{fc_interface.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical performance metrics for a Fibre Channel interface for a specific time.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fc_interface.uuid | string | path | True | The unique identifier of the Fibre Channel interface. |
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the FC interface. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
|---|---|---|
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an FC interface

DELETE `/network/fc/interfaces/{uuid}`

**Introduced In:** 9.6

Deletes an FC interface.

#### Related ONTAP commands

- `network interface delete`

#### Learn more

- DOC /network/fc/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53280992 | The FC interface could not be deleted because it is enabled. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC interface

GET `/network/fc/interfaces/{uuid}`

**Introduced In:** 9.6

Retrieves an FC interface.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `network interface show`
- `vserver fcp interface show`

#### Learn more

- DOC /network/fc/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier for the FC interface. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| data_protocol | string | The data protocol for which the FC interface is configured. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |

| Name | Type | Description |
| --- | --- | --- |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "data_protocol": "string",
  "location": {
    "home_node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
      },
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "fc_lif1",
  "port_address": "5060F",
  "state": "string",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
```

```
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "bce9827d-4d8f-60af-c771-6e8e9af2c6f0",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

home_node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

home_port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.

`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship.

| Name | Type | Description |
| --- | --- | --- |
| home_node | home_node | |
| home_port | home_port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| is_home | boolean | Indicates if the FC interface is currently on its home node. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an FC interface

PATCH `/network/fc/interfaces/{uuid}`

**Introduced In:** 9.6

Updates an FC interface.

#### Related ONTAP commands

- `network interface modify`

#### Learn more

- DOC /network/fc/interfaces

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier for the FC interface. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br><br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br><br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

### Example request


```
{
  "comment": "string",
  "location": {
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "fc_lif1",
  "port_address": "5060F",
  "state": "string",
  "uuid": "bce9827d-4d8f-60af-c771-6e8e9af2c6f0",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
```


### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966140 | An interface with the same name already exists. |
| 1966217 | The specified port is not valid on the node provided. |
| 1966238 | The node or port of an active SAN data interface cannot be changed. |
| 1966702 | The destination node is not healthy. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5374579 | The SAN Kernel Agent on the node is unavailable. |
| 5374870 | A partial failure occurred; renaming the interface failed. Correct the error and resubmit the request. |
| 5374871 | The Fibre Channel port identified by the specified UUID does not refer to the same port as that identified by the specified node name and/or port name. |
| 5374872 | If either `location.port.node.name` or `location.port.name` is supplied, both properties must be supplied. |
| 5375057 | An FC port with the provided UUID does not exist. |
| 72089674 | You cannot move a Fibre Channel interface configured for the NVMe over FC data protocol. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

home_node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

home_port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC

network.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.

`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship.

| Name | Type | Description |
|---|---|---|
| home_node | home_node | |
| home_port | home_port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |
| is_home | boolean | Indicates if the FC interface is currently on its home node. |
| node | node | |
| port | port | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fc_interface

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the cluster node and port names or the port UUID.

| Name | Type | Description |
|---|---|---|
| comment | string | A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH. |
| enabled | boolean | The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to *true* (enabled) in POST. |
| location | location | The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its cluster node name and port name. Either the UUID or the cluster node name and port name are required for POST. To move an interface, supply either the port UUID or the cluster node name and port name in a PATCH.<br><br>`location.node` and `location.port` refer to the current location of the FC interface. This can be different from `location.home_node` and `location.home_port` in instances where the FC interface has failed over to its HA partner node. The `location.node`, `location.port`, and `location.is_home` properties are not available for interfaces on the inactive side of a MetroCluster relationship. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the FC interface. Required in POST; optional in PATCH. |
| port_address | string | The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| state | string | The current operational state of the FC interface. The state is set to *down* if the interface is not enabled.<br><br>If the cluster node hosting the port is down or unavailable, no state value is returned. |
| uuid | string | The unique identifier of the FC interface. Required in the URL. |
| wwnn | string | The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.<br><br>• example: 20:00:00:50:56:b4:13:01<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| wwpn | string | The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.<br>• example: 20:00:00:50:56:b4:13:a8<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve FC port information

### Retrieve FC port information

#### Overview

Fibre Channel (FC) ports are the physical ports of FC adapters on ONTAP cluster nodes that can be connected to FC networks to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

The Fibre Channel port REST API allows you to discover FC ports, obtain status information for FC ports, and configure FC port properties. POST and DELETE requests are not supported. You must physically add and remove FC adapters to ONTAP nodes to create and remove ports from the ONTAP cluster.

#### Performance monitoring

Performance of an FC port can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the performance of an FC port in terms of IOPS, latency, and throughput. The `metric.*` properties denote an average, whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

### Retrieving all FC ports

```
# The API:
GET /api/network/fc/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports" -H "Accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
    "name": "0a",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
```

```
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
    "name": "0b",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
    "name": "0c",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2748-b047-11e8-9af3-005056bb838e",
```

```
    "name": "0d",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2748-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b28c2-b047-11e8-9af3-005056bb838e",
    "name": "0e",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b28c2-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2a7b-b047-11e8-9af3-005056bb838e",
    "name": "0f",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2a7b-b047-11e8-9af3-
005056bb838e"
      }
```

```
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2e2b-b047-11e8-9af3-005056bb838e",
    "name": "1b",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2e2b-b047-11e8-9af3-
005056bb838e"
      }
    }
  }
],
"num_records": 8,
"_links": {
  "self": {
    "href": "/api/network/fc/ports"
  }
}
}
```


#### Retrieving all FC ports with state *online*

The `state` query parameter is used to perform the query.

```
# The API:
GET /api/network/fc/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports?state=online" -H
"Accept: application/hal+json"

# The response:
{
```

```
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
    "name": "0a",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
    "name": "0b",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
```

```
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
    "name": "0c",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-
005056bb838e"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/network/fc/ports?state=online"
  }
}
}
```


---

### Retrieving an FC port

```
# The API:
GET /api/network/fc/ports/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-
9af3-005056bb838e" -H "Accept: application/hal+json"

# The response:
{
"node": {
  "name": "node1",
  "uuid": "5a534a72-b047-11e8-9af3-005056bb838e",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/5a534a72-b047-11e8-9af3-005056bb838e"
```

```
      }
    }
  },
  "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
  "name": "0a",
  "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
  Adapter, rev. 1.0.0, 8G)",
  "enabled": true,
  "fabric": {
    "connected": true,
    "connected_speed": 8,
    "name": "55:0e:b1:a0:20:40:80:00",
    "port_address": "52100",
    "switch_port": "ssan-g620-03:1"
  },
  "physical_protocol": "fibre_channel",
  "speed": {
    "maximum": "8",
    "configured": "auto"
  },
  "state": "online",
  "supported_protocols": [
    "fcp"
  ],
  "transceiver": {
    "form_factor": "SFP",
    "manufacturer": "ACME",
    "capabilities": [
      4,
      8
    ],
    "part_number": "1000"
  },
  "wwnn": "50:0a:09:80:bb:83:8e:00",
  "wwpn": "50:0a:09:82:bb:83:8e:00",
  "metric": {
    "timestamp": "2019-04-09T05:50:15Z",
    "duration": "PT15S",
    "status": "ok",
    "latency": {
      "other": 0,
      "total": 0,
      "read": 0,
      "write": 0
    },
    "iops": {
```

```
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-005056bb838e"
  }
}
}
```


### Disabling an FC port

If an active FC interface exists on an FC port, the port cannot be disabled.

```
# The API:
PATCH /api/network/fc/ports/{uuid}

# The call:
curl -X PATCH "http://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-
9af3-005056bb838e" -H "Accept: application/hal+json" -d '{ "enabled":
false }'
```

### Retrieve FC ports

GET `/network/fc/ports`

**Introduced In:** 9.6

Retrieves FC ports.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `fabric.name`
- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `network fcp adapter show`

#### Learn more

- DOC /network/fc/ports

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| transceiver.capabilities | integer | query | False | Filter by transceiver.capabilities |
| transceiver.form_factor | string | query | False | Filter by transceiver.form_factor<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| transceiver.part_number | string | query | False | Filter by transceiver.part_number |
| transceiver.manufacturer | string | query | False | Filter by transceiver.manufacturer |
| fabric.connected_speed | integer | query | False | Filter by fabric.connected_speed |
| fabric.switch_port | string | query | False | Filter by fabric.switch_port |
| fabric.connected | boolean | query | False | Filter by fabric.connected |
| fabric.name | string | query | False | Filter by fabric.name |
| fabric.port_address | string | query | False | Filter by fabric.port_address |
| supported_protocols | string | query | False | Filter by supported_protocols |
| physical_protocol | string | query | False | Filter by physical_protocol |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br>• Introduced in: 9.8 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br>• Introduced in: 9.8 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.8 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.8 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.8 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.8 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.8 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.8 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.8 |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.8 |
| interface_count | integer | query | False | Filter by interface_count<br><br>• Introduced in: 9.10 |
| uuid | string | query | False | Filter by uuid |
| wwnn | string | query | False | Filter by wwnn |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.8 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br>• Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.8 |
| description | string | query | False | Filter by description |
| speed.maximum | string | query | False | Filter by speed.maximum |
| speed.configured | string | query | False | Filter by speed.configured |
| name | string | query | False | Filter by name |
| state | string | query | False | Filter by state |
| enabled | boolean | query | False | Filter by enabled |
| wwpn | string | query | False | Filter by wwpn |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fc_port] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "Fibre Channel Target Adapter 0a (ACME Fibre
Channel Adapter, rev. 1.0.0, 8G)",
      "fabric": {
        "connected_speed": 16,
        "name": "string",
        "port_address": "52100A",
        "switch_port": "ssan-g620-03:33"
      },
      "interface_count": 0,
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
```

```
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "0a",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "physical_protocol": "string",
      "speed": {
        "configured": "auto",
        "maximum": "32"
      },
      "state": "online",
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "supported_protocols": [
        "string"
      ],
      "transceiver": {
```

```
        "capabilities": [
          16
        ],
        "form_factor": "string",
        "manufacturer": "Acme, Inc.",
        "part_number": "string"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "wwnn": "20:00:00:50:56:b4:13:a8",
      "wwpn": "20:00:00:50:56:b4:13:a8"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fabric

Properties of the fabric to which the FC port is attached.

| Name | Type | Description |
|---|---|---|
| connected | boolean | Reports if the physical port has established a connection with the FC fabric. |
| connected_speed | integer | The negotiated data rate between the target FC port and the fabric in gigabits per second. |
| name | string | The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The FC port address of the host bus adapter (HBA) physical port.<br><br>Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a six-digit hexadecimal encoded numeric value. |
| switch_port | string | The switch port to which the FC port is connected. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

speed

The physical device speed related properties of the FC port.

| Name | Type | Description |
| --- | --- | --- |
| configured | string | The configured speed of the FC port in gigabits per second. |
| maximum | string | The maximum speed supported by the FC port in gigabits per second. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

transceiver

Properties of the transceiver connected to the FC port.

| Name | Type | Description |
| --- | --- | --- |
| capabilities | array[integer] | The speeds of which the transceiver is capable in gigabits per second. |

| Name | Type | Description |
| --- | --- | --- |
| form_factor | string | The form factor of the transceiver. Possible values are:<br>• *sfp* - Small Form Factor - Pluggable<br>• *sff* - Small Form Factor<br>• *unknown* - Unknown |
| manufacturer | string | The manufacturer of the transceiver. |
| part_number | string | The part number of the transceiver. |

fc_port

A Fibre Channel (FC) port is the physical port of an FC adapter on an ONTAP cluster node that can be connected to an FC network to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| description | string | A description of the FC port. |
| enabled | boolean | The administrative state of the FC port. If this property is set to *false*, all FC connectivity to FC interfaces are blocked. Optional in PATCH. |
| fabric | fabric | Properties of the fabric to which the FC port is attached. |
| interface_count | integer | The number of FC interfaces currently provisioned on this port. This property is not supported in an SVM context. |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The FC port name. |
| node | node | |

| Name | Type | Description |
| --- | --- | --- |
| physical_protocol | string | The physical network protocol of the FC port. |
| speed | speed | The physical device speed related properties of the FC port. |
| state | string | The operational state of the FC port.<br>• startup - The port is booting up.<br>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.<br>• online - The port is initialized and a link with the fabric has been established.<br>• link_disconnected - The link was present at one point on this port but is currently not established.<br>• offlined_by_user - The port is administratively disabled.<br>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.<br>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| supported_protocols | array[string] | The network protocols supported by the FC port. |
| transceiver | transceiver | Properties of the transceiver connected to the FC port. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the FC port. |
| wwnn | string | The base world wide node name (WWNN) for the FC port. |
| wwpn | string | The base world wide port name (WWPN) for the FC port. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for an FC port**

GET `/network/fc/ports/{fc_port.uuid}/metrics`

**Introduced In:** 9.14

Retrieves historical performance metrics for a Fibre Channel port.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| status | string | query | False | Filter by status |
| uuid | string | query | False | Filter by uuid |
| duration | string | query | False | Filter by duration |
| fc_port.uuid | string | path | True | The unique identifier of the Fibre Channel port. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the FC port. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve historical performance metrics for an FC port for a specific time

GET `/network/fc/ports/{fc_port.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical performance metrics for a Fibre Channel port for a specific time.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fc_port.uuid | string | path | True | The unique identifier of the Fibre Channel port. |
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the FC port. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
|---|---|---|
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC port

GET `/network/fc/ports/{uuid}`

**Introduced In:** 9.6

Retrieves an FC port.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `fabric.name`
- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `network fcp adapter show`

#### Learn more

- DOC /network/fc/ports

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier for the FC port. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| description | string | A description of the FC port. |
| enabled | boolean | The administrative state of the FC port. If this property is set to *false*, all FC connectivity to FC interfaces are blocked. Optional in PATCH. |
| fabric | fabric | Properties of the fabric to which the FC port is attached. |

| Name | Type | Description |
| --- | --- | --- |
| interface_count | integer | The number of FC interfaces currently provisioned on this port. This property is not supported in an SVM context. |
| metric | metric | Performance numbers, such as IOPS latency and throughput |
| name | string | The FC port name. |
| node | node | |
| physical_protocol | string | The physical network protocol of the FC port. |
| speed | speed | The physical device speed related properties of the FC port. |
| state | string | The operational state of the FC port.<br>• startup - The port is booting up.<br>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.<br>• online - The port is initialized and a link with the fabric has been established.<br>• link_disconnected - The link was present at one point on this port but is currently not established.<br>• offlined_by_user - The port is administratively disabled.<br>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.<br>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| supported_protocols | array[string] | The network protocols supported by the FC port. |
| transceiver | transceiver | Properties of the transceiver connected to the FC port. |
| uuid | string | The unique identifier of the FC port. |
| wwnn | string | The base world wide node name (WWNN) for the FC port. |
| wwpn | string | The base world wide port name (WWPN) for the FC port. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
Adapter, rev. 1.0.0, 8G)",
  "fabric": {
    "connected_speed": 16,
    "name": "string",
    "port_address": "52100A",
    "switch_port": "ssan-g620-03:33"
  },
  "interface_count": 0,
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "0a",
  "node": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "string",
  "speed": {
    "configured": "auto",
    "maximum": "32"
  },
  "state": "online",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "supported_protocols": [
    "string"
  ],
  "transceiver": {
    "capabilities": [
      16
    ],
    "form_factor": "string",
    "manufacturer": "Acme, Inc.",
    "part_number": "string"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:a8",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fabric

Properties of the fabric to which the FC port is attached.

| Name | Type | Description |
|---|---|---|
| connected | boolean | Reports if the physical port has established a connection with the FC fabric. |
| connected_speed | integer | The negotiated data rate between the target FC port and the fabric in gigabits per second. |
| name | string | The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The FC port address of the host bus adapter (HBA) physical port.<br><br>Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a six-digit hexadecimal encoded numeric value. |
| switch_port | string | The switch port to which the FC port is connected. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

speed

The physical device speed related properties of the FC port.

| Name | Type | Description |
| --- | --- | --- |
| configured | string | The configured speed of the FC port in gigabits per second. |
| maximum | string | The maximum speed supported by the FC port in gigabits per second. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

transceiver

Properties of the transceiver connected to the FC port.

| Name | Type | Description |
| --- | --- | --- |
| capabilities | array[integer] | The speeds of which the transceiver is capable in gigabits per second. |

| Name | Type | Description |
|---|---|---|
| form_factor | string | The form factor of the transceiver. Possible values are:<br>• *sfp* - Small Form Factor - Pluggable<br>• *sff* - Small Form Factor<br>• *unknown* - Unknown |
| manufacturer | string | The manufacturer of the transceiver. |
| part_number | string | The part number of the transceiver. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update an FC port**

PATCH `/network/fc/ports/{uuid}`

**Introduced In:** 9.6

Updates an FC port.

#### Related ONTAP commands

- `network fcp adapter modify`

#### Learn more

- DOC /network/fc/ports

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier for the FC port. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| description | string | A description of the FC port. |
| enabled | boolean | The administrative state of the FC port. If this property is set to *false*, all FC connectivity to FC interfaces are blocked. Optional in PATCH. |
| fabric | fabric | Properties of the fabric to which the FC port is attached. |
| interface_count | integer | The number of FC interfaces currently provisioned on this port. This property is not supported in an SVM context. |
| name | string | The FC port name. |
| node | node | |
| physical_protocol | string | The physical network protocol of the FC port. |
| speed | speed | The physical device speed related properties of the FC port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The operational state of the FC port.<br>• startup - The port is booting up.<br>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.<br>• online - The port is initialized and a link with the fabric has been established.<br>• link_disconnected - The link was present at one point on this port but is currently not established.<br>• offlined_by_user - The port is administratively disabled.<br>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.<br>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible. |
| supported_protocols | array[string] | The network protocols supported by the FC port. |
| transceiver | transceiver | Properties of the transceiver connected to the FC port. |
| uuid | string | The unique identifier of the FC port. |
| wwnn | string | The base world wide node name (WWNN) for the FC port. |
| wwpn | string | The base world wide port name (WWPN) for the FC port. |

#### Example request


```
{
  "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
Adapter, rev. 1.0.0, 8G)",
  "fabric": {
    "connected_speed": 16,
    "name": "string",
    "port_address": "52100A",
    "switch_port": "ssan-g620-03:33"
  },
  "interface_count": 0,
  "name": "0a",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "string",
  "speed": {
    "configured": "auto",
    "maximum": "32"
  },
  "state": "online",
  "supported_protocols": [
    "string"
  ],
  "transceiver": {
    "capabilities": [
      16
    ],
    "form_factor": "string",
    "manufacturer": "Acme, Inc.",
    "part_number": "string"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:a8",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
```


#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5374085 | The node where the Fibre Channel port is located is offline. |
| 5374087 | The Fibre Channel port has active Fibre Channel interfaces and cannot be disabled. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

fabric

Properties of the fabric to which the FC port is attached.

| Name | Type | Description |
|---|---|---|
| connected | boolean | Reports if the physical port has established a connection with the FC fabric. |
| connected_speed | integer | The negotiated data rate between the target FC port and the fabric in gigabits per second. |
| name | string | The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The FC port address of the host bus adapter (HBA) physical port.<br><br>Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a six-digit hexadecimal encoded numeric value. |
| switch_port | string | The switch port to which the FC port is connected. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

speed

The physical device speed related properties of the FC port.

| Name | Type | Description |
| --- | --- | --- |
| configured | string | The configured speed of the FC port in gigabits per second. |
| maximum | string | The maximum speed supported by the FC port in gigabits per second. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

transceiver

Properties of the transceiver connected to the FC port.

| Name | Type | Description |
| --- | --- | --- |
| capabilities | array[integer] | The speeds of which the transceiver is capable in gigabits per second. |

| Name | Type | Description |
| --- | --- | --- |
| form_factor | string | The form factor of the transceiver. Possible values are:<br>• *sfp* - Small Form Factor - Pluggable<br>• *sff* - Small Form Factor<br>• *unknown* - Unknown |
| manufacturer | string | The manufacturer of the transceiver. |
| part_number | string | The part number of the transceiver. |

fc_port

A Fibre Channel (FC) port is the physical port of an FC adapter on an ONTAP cluster node that can be connected to an FC network to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

| Name | Type | Description |
| --- | --- | --- |
| description | string | A description of the FC port. |
| enabled | boolean | The administrative state of the FC port. If this property is set to *false*, all FC connectivity to FC interfaces are blocked. Optional in PATCH. |
| fabric | fabric | Properties of the fabric to which the FC port is attached. |
| interface_count | integer | The number of FC interfaces currently provisioned on this port. This property is not supported in an SVM context. |
| name | string | The FC port name. |
| node | node | |
| physical_protocol | string | The physical network protocol of the FC port. |
| speed | speed | The physical device speed related properties of the FC port. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The operational state of the FC port.<br>• startup - The port is booting up.<br>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.<br>• online - The port is initialized and a link with the fabric has been established.<br>• link_disconnected - The link was present at one point on this port but is currently not established.<br>• offlined_by_user - The port is administratively disabled.<br>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.<br>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible. |
| supported_protocols | array[string] | The network protocols supported by the FC port. |
| transceiver | transceiver | Properties of the transceiver connected to the FC port. |
| uuid | string | The unique identifier of the FC port. |
| wwnn | string | The base world wide node name (WWNN) for the FC port. |
| wwpn | string | The base world wide port name (WWPN) for the FC port. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage HTTP proxy configuration

### Manage HTTP proxy configuration

#### Overview

Configuration of an HTTP proxy for an SVM or a Cluster IPspace.

#### Retrieve HTTP proxy information

The HTTP proxy GET operation retrieves all configurations for an SVM or a Cluster IPspace via '/api/cluster'.

#### Examples

### Retrieving all fields for all HTTP proxy configurations

```
# The API:
/api/network/http-proxy

# The call:
curl -X GET "https://<mgmt-ip>/api/network/http-
proxy?fields=*&return_records=true&return_timeout=15" -H  "accept:
application/json"

# The response:
{
"records": [
  {
    "uuid": "4133a1fc-7228-11e9-b40c-005056bb4f0c",
    "svm": {
        "name": "vs1",
        "uuid": "4133a1fc-7228-11e9-b40c-005056bb4f0c"
    },
    "server": "server1.example.com",
    "port": 3128,
    "authentication_enabled": false
  },
  {
    "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c",
    "svm": {
        "name": "cluster-1",
        "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c"
    },
    "ipspace": {
        "uuid": "7433520f-7214-11e9-828c-005056bb4f0c",
        "name": "Default"
    },
    "server": "1.1.1.",
    "port": 3128,
    "authentication_enabled": true
   }
],
"num_records": 2
}
```

**Retrieving the HTTP proxy configuration for a specific SVM**

```
# The API:
/api/network/http-proxy/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-11e9-
828c-005056bb4f0c" -H  "accept: application/json"

# The response
{
"uuid": "96219ce3-7214-11e9-828c-005056bb4f0c",
"svm": {
    "name": "cluster-1",
    "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c"
},
"ipspace": {
    "uuid": "7433520f-7214-11e9-828c-005056bb4f0c",
    "name": "Default"
},
"server": "1.1.1.1",
"port": 3128,
"authentication_enabled": false
}
```

#### Creating an HTTP proxy configuration

You can use the HTTP proxy POST operation to create an HTTP proxy configuration for the specified SVM.

#### Examples

### Creating an HTTP proxy configuration for a particular SVM

```
# The API:
/api/network/http-proxy

# The call:
curl -X POST "https://<mgmt-ip>/api/network/http-proxy" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{    \"port\":
3128,  \"server\": \"1.1.1.1\",  \"svm\": {      \"name\": \"cluster-1\"
}}"
```

### Creating an HTTP proxy configuration for a particular IPspace

```
# The API:
/api/network/http-proxy

# The call:
curl -X POST "https://<mgmt-ip>/api/network/http-proxy" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{  \"ipspace\":
{    \"name\": \"Default\"  },  \"port\": 3128,  \"server\": \"1.1.1.1\"}"
```

### Creating an HTTP proxy configuration with authentication enabled

```
# The API:
/api/network/http-proxy

# The call:
curl -X POST "https://<mgmt-ip>/api/network/http-proxy" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{  \"ipspace\":
{    \"name\": \"Default\"  },  \"port\": 3128,  \"server\": \"1.1.1.1\",
\"authentication_enabled\":true, \"username\":\"test\",
\"password\":\"test\"}"
```

#### Update an HTTP proxy configuration for a specified SVM

You can use the HTTP proxy PATCH operation to update the HTTP proxy configuration for the specified SVM.

#### Example

The following example shows how a PATCH operation is used to update an HTTP proxy configuration for a specific SVM:

```
# The API:
/api/network/http-proxy/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-
11e9-828c-005056bb4f0c" -H  "accept: application/json" -H  "Content-Type:
application/json" -d "{    \"port\": 3128,  \"server\":
\"server2.example.com\"}"
```

#### Delete an HTTP proxy configuration for a specified SVM

You can use the HTTP proxy DELETE operation to delete the HTTP proxy configuration for the specified SVM.

#### Example

The following example shows how a DELETE operation is used to delete an HTTP proxy configuration for a specific SVM:

```
# The API:
/api/network/http-proxy/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-
11e9-828c-005056bb4f0c" -H  "accept: application/json"
```

### Retrieve HTTP proxy configurations for all SVMs and cluster IPspaces

GET `/network/http-proxy`

**Introduced In:** 9.7

Retrieves the HTTP proxy configurations of all the SVMs and Cluster IPspaces.

#### Related ONTAP commands

- `vserver http-proxy show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| authentication_enabled | boolean | query | False | Filter by authentication_enabled<br>• Introduced in: 9.9 |
| server | string | query | False | Filter by server |
| scope | string | query | False | Filter by scope |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| uuid | string | query | False | Filter by uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| port | integer | query | False | Filter by port<br>• Max value: 65535<br>• Min value: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of HTTP proxy records |
| records | array[network_http_proxy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": 3128,
      "scope": "string",
      "server": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This is mutually exclusive with IPspace during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

network_http_proxy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH. |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |
| svm | svm | This is mutually exclusive with IPspace during POST and PATCH. |
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an HTTP proxy configuration for an SVM or cluster IPspace

POST `/network/http-proxy`

**Introduced In:** 9.7

Creates an HTTP proxy configuration for an SVM or a Cluster IPspace.
Important notes:

- IPv6 must be enabled if IPv6 family addresses are specified in the "server" field.
- The server and the port combination specified using the "server" and "port" fields is validated during this operation. The validation will fail in the following scenarios:
  - The HTTP proxy service is not configured on the server.
  - The HTTP proxy service is not running on the specified port.
  - The server is unreachable.

#### Required properties

- SVM-scoped HTTP proxy
  - `svm.uuid` or `svm.name` - Existing SVM in which to create the HTTP proxy.
- Cluster-scoped HTTP proxy
  - `ipspace.uuid` or `ipspace.name` - Existing Cluster IPspace in which to create the HTTP proxy.
- `server` - HTTP proxy server FQDN or IP address.
- `port` - HTTP proxy server port.

#### Optional properties

- `authentication_enabled` - Specifies if authentication is required for the HTTP proxy server.
- `username` - Username used to authenticate with the HTTP proxy server.
- `password` - Password used to authenticate with the HTTP proxy server.

#### Related ONTAP commands

- `vserver http-proxy create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH. |
| password | string | Password to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |
| svm | svm | This is mutually exclusive with IPspace during POST and PATCH. |
| username | string | Username to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

#### Example request

```
{
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "password": "string",
  "port": 3128,
  "scope": "string",
  "server": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "username": "string",
  "uuid": "string"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of HTTP proxy records |
| records | array[network_http_proxy] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "password": "string",
      "port": 3128,
      "scope": "string",
      "server": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "username": "string",
      "uuid": "string"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 26214473 | HTTP proxy configuration is not valid. |
| 26214476 | The "IPspace" parameter should not be specified in the SVM context. |
| 26214477 | The specified IPspace does not exist. |

| Error Code | Description |
| --- | --- |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces. |
| 26214481 | Username and password cannot be empty when HTTP proxy authentication is enabled. |
| 26214482 | Username and password cannot be specified when HTTP proxy authentication is disabled. |
| 26214480 | One of "svm.name", "svm.uuid", "ipspace.name" or "ipspace.uuid" must be specified. |
| 2621462 | SVM "vs0" does not exist. |
| 262186 | Field "svm.name" cannot be used with field "ipspace.name". |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This is mutually exclusive with IPspace during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

network_http_proxy

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| password | string | Password to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |
| svm | svm | This is mutually exclusive with IPspace during POST and PATCH. |
| username | string | Username to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete HTTP proxy configuration for an SVM or cluster IPspace

DELETE `/network/http-proxy/{uuid}`

**Introduced In:** 9.7

Deletes the HTTP proxy configuration of the specified SVM or Cluster IPspace.

#### Related ONTAP commands

- `vserver http-proxy delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | HTTP proxy UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 138281013 | The HTTP Proxy cannot be deleted while in use by a cloud agent connection. |
| 2621574 | This operation is not permitted on a SVM that is configured as the destination of a Metrocluster SVM relationship. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Display HTTP proxy server, port, and IPspace information for an SVM or cluster IPspace

GET `/network/http-proxy/{uuid}`

**Introduced In:** 9.7

Displays the HTTP proxy server, port, and IPspace of the specified SVM or Cluster IPspace.

#### Related ONTAP commands

- `vserver http-proxy show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | HTTP proxy UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH. |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | This is mutually exclusive with IPspace during POST and PATCH. |
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": 3128,
  "scope": "string",
  "server": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This is mutually exclusive with IPspace during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update proxy server, port, username, and password parameters

PATCH `/network/http-proxy/{uuid}`

**Introduced In:** 9.7

Updates the proxy server, port, username, and password parameters.
Important notes:

- IPv6 must be enabled if IPv6 family addresses are specified in the "server" field.
- The server and the port combination specified using the "server" and "port" fields is validated during this operation. The validation will fail in the following scenarios:
  - The HTTP proxy service is not configured on the server.
  - The HTTP proxy service is not running on the specified port.
  - The server is unreachable.

#### Related ONTAP commands

- `vserver http-proxy modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | HTTP proxy UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| password | string | Password to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |
| username | string | Username to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

**Example request**

```
{
  "password": "string",
  "port": 3128,
  "scope": "string",
  "server": "string",
  "username": "string",
  "uuid": "string"
}
```

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 26214473 | The HTTP proxy configuration is not valid. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces. |
| 26214481 | Username and password cannot be empty when HTTP proxy authentication is enabled. |
| 26214482 | Username and password cannot be specified when HTTP proxy authentication is disabled. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input. This is mutually exclusive with SVM during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

This is mutually exclusive with IPspace during POST and PATCH.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

network_http_proxy

| Name | Type | Description |
|---|---|---|
| authentication_enabled | boolean | Specifies whether or not authentication with the HTTP proxy server is enabled. |
| password | string | Password to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |

| Name | Type | Description |
| --- | --- | --- |
| port | integer | The port number on which the HTTP proxy service is configured on the<br>proxy server. |
| scope | string | Set to “svm” for HTTP proxy owned by an SVM. Otherwise, set to "cluster". |
| server | string | Fully qualified domain name (FQDN) or IP address of the HTTP proxy server. |
| username | string | Username to authenticate with the HTTP proxy server when authentication_enabled is set to "true". |
| uuid | string | The UUID that uniquely identifies the HTTP proxy. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage BGP peer groups

## Manage BGP peer groups

### Overview

The following operations are supported:

- Creation: POST network/ip/bgp/peer-groups
- Collection Get: GET network/ip/bgp/peer-groups
- Instance Get: GET network/ip/bgp/peer-groups/{uuid}
- Instance Patch: PATCH network/ip/bgp/peer-groups/{uuid}
- Instance Delete: DELETE network/ip/bgp/peer-groups/{uuid}

### Retrieving network BGP sessions information

The IP BGP peer-groups GET API retrieves and displays relevant information pertaining to the BGP peer-groups configured in the cluster. The response can contain a list of multiple BGP peer-groups or a specific peer-group. Each BGP peer-group represents a BGP session configured between a local interface and a peer router.

### Examples

#### Retrieving all BGP peer-groups in the cluster

The following example shows the list of all BGP peer-groups configured in a cluster.

```
# The API:
/api/network/ip/bgp/peer-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-groups" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "5f22ae9d-87b2-11e9-a3a6-005056bb81a4",
    "name": "pg1",
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/5f22ae9d-87b2-11e9-a3a6-
005056bb81a4"
      }
    }
  },
  {
    "uuid": "5fd08be3-87b2-11e9-952f-005056bb2170",
    "name": "pg2",
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/5fd08be3-87b2-11e9-952f-
005056bb2170"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/ip/bgp/peer-groups"
  }
}
}
```

#### Retrieving a specific BGP peer-group

The following example shows the response when a specific BGP peer-group is requested. The system returns an error when there is no peer-group with the requested UUID.

```
# The API:
/api/network/ip/bgp/peer-groups/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/5fd08be3-
87b2-11e9-952f-005056bb2170" -H "accept: application/hal+json"

# The response:
{
"uuid": "5fd08be3-87b2-11e9-952f-005056bb2170",
"name": "pg2",
"ipspace": {
  "uuid": "84fd3375-879a-11e9-a3a6-005056bb81a4",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/84fd3375-879a-11e9-a3a6-005056bb81a4"
    }
  }
},
"local": {
  "interface": {
    "uuid": "5e76a305-87b2-11e9-952f-005056bb2170",
    "name": "bgp2",
    "ip": {
      "address": "10.10.10.2"
    }
  },
  "port": {
    "uuid": "f8ff73de-879a-11e9-952f-005056bb2170",
    "name": "e0h",
    "node": {
      "name": "node1"
    }
  }
},
"peer": {
  "address": "10.10.10.1",
  "asn": 65501
},
"state": "up",
"_links": {
  "self": {
    "href": "/api/network/ip/bgp/peer-groups/5fd08be3-87b2-11e9-952f-
005056bb2170"
```

```
  }
}
}
```

#### Retrieving specific fields and limiting the output using filters

The following example shows the response when a filter is applied (location.port.node.name=node1) and only certain fields are requested. Filtered fields are in the output in addition to the default fields and requested fields.

```
# The API:
/api/network/ip/bgp/peer-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-
groups?local.port.node.name=node1&fields=local.interface.ip,peer" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "5f22ae9d-87b2-11e9-a3a6-005056bb81a4",
    "name": "pg1",
    "local": {
      "interface": {
        "ip": {
          "address": "10.10.10.1"
        }
      },
      "port": {
        "node": {
          "name": "node1"
        }
      }
    },
    "peer": {
      "address": "10.10.10.2",
      "asn": 65501
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/5f22ae9d-87b2-11e9-a3a6-
```

```
005056bb81a4"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/bgp/peer-
groups?local.port.node.name=node1&fields=local.interface.ip,peer"
  }
}
}
```


#### Creating a BGP peer-group

The BGP peer-group POST API is used to create a peer-group as shown in the following examples.

#### Examples

### Creating a BGP peer-group with an existing interface

The following example shows how to create a BGP peer-group between an existing interface "bgp1" and peer router with the address "10.10.10.10". The local interface "bgp1" needs to support the management-bgp service, otherwise the system returns an error.

```
# The API:
/api/network/ip/bgp/peer-groups

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/bgp/peer-
groups?return_records=true" -d'{"name": "newPg", "ipspace.name":"Default",
"local.interface.name": "bgp1", "peer.address":"10.10.10.10"}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "e3faacc6-87cb-11e9-a3a6-005056bb81a4",
    "name": "newPg",
    "ipspace": {
      "name": "Default"
    },
    "local": {
      "interface": {
        "name": "bgp1"
      }
    },
    "peer": {
      "address": "10.10.10.10"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/e3faacc6-87cb-11e9-a3a6-
005056bb81a4"
      }
    }
  }
]
}
```

### Creating a BGP peer-group and provisioning a new local interface

The following example shows how to create a BGP peer-group with any local interface. If the local interface doesn’t exist, the system will create it first before creating the peer-group.

```
# The API:
/api/network/ip/bgp/peer-groups

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/bgp/peer-
groups?return_records=true" -d'{"name": "newPg1",
"ipspace.name":"Default", "local": {"interface": {"name": "newlif"}, "ip":
{"address": "9.9.9.9", "netmask": "24"}, "port": {"name": "e0f", "node":
{"name": "node1"}}}, "peer.address":"10.10.10.10"}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "c292f069-8872-11e9-a3a6-005056bb81a4",
    "name": "newPg1",
    "ipspace": {
      "name": "Default"
    },
    "local": {
      "interface": {
        "name": "newlif"
      },
      "port": {
        "name": "e0f",
        "node": {
          "name": "node1"
        }
      }
    },
    "peer": {
      "address": "10.10.10.10"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/c292f069-8872-11e9-a3a6-
005056bb81a4"
      }
    }
  }
]
}
```

#### Updating BGP peer-groups

The BGP peer-groups PATCH API is used to update attributes of a peer-group.

#### Examples

### Updating the peer router address

The following example shows how the PATCH request changes the peer router IP address.

```
# The API:
/api/network/ip/bgp/peer-groups/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-
1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d
'{"peer.address": "10.10.10.20" }'
{
}
```

### Updating the peer-group to a new name

The following example shows how the PATCH request renames the peer-group.

```
# The API:
/api/network/ip/bgp/peer-groups/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-
1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d
'{"name": "NewName"}'
{
}
```

#### Deleting BGP peer-groups

The BGP peer-groups DELETE API is used to delete an BGP peer-group.

#### Example

### Deleting a BGP peer-group

The following DELETE request deletes a BGP peer-group.

```
# The API:
/api/network/ip/bgp/peer-group/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-
1f43-11e9-803e-005056a7646a"
{
}
```

### Retrieve all BGP peer group details for VIP

GET `/network/ip/bgp/peer-groups`

**Introduced In:** 9.7

Retrieves the details of all BGP peer groups for VIP.

#### Related ONTAP Commands

- `network bgp peer-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| peer.asn | integer | query | False | Filter by peer.asn |
| peer.is_next_hop | boolean | query | False | Filter by peer.is_next_hop<br>• Introduced in: 9.9 |
| peer.address | string | query | False | Filter by peer.address |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| peer.md5_enabled | boolean | query | False | Filter by peer.md5_enabled<br>• Introduced in: 9.16 |
| peer.md5_secret | string | query | False | Filter by peer.md5_secret<br>• Introduced in: 9.16<br>• maxLength: 79<br>• minLength: 1 |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| uuid | string | query | False | Filter by uuid |
| local.interface.ip.address | string | query | False | Filter by local.interface.ip.address |
| local.interface.name | string | query | False | Filter by local.interface.name |
| local.interface.uuid | string | query | False | Filter by local.interface.uuid |
| local.port.node.name | string | query | False | Filter by local.port.node.name |
| local.port.uuid | string | query | False | Filter by local.port.uuid |
| local.port.name | string | query | False | Filter by local.port.name |
| name | string | query | False | Filter by name |
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[bgp_peer_group] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "local": {
        "interface": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "e1b",
          "node": {
            "name": "node1"
          },
```

```
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "name": "bgpv4peer",
        "peer": {
          "address": "10.10.10.7",
          "md5_enabled": 1,
          "md5_secret": "SECRET_WORD"
        },
        "state": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ip

IP information to create a new interface.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

local

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

peer

Information describing the router to peer with

| Name | Type | Description |
| --- | --- | --- |
| address | string | Peer router address |
| asn | integer | Autonomous system number of peer |
| is_next_hop | boolean | Use peer address as next hop. |
| md5_enabled | boolean | Enable or disable TCP MD5 signatures per RFC 2385. |
| md5_secret | string | The shared TCP MD5 secret key. This can either be given as a password or hexadecimal key. |

bgp_peer_group

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Either the UUID or name is supplied on input. |
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |

| Name | Type | Description |
|---|---|---|
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new BGP peer group for VIP**

POST `/network/ip/bgp/peer-groups`

**Introduced In:** 9.7

Creates a new BGP peer group for VIP. Multipath-routing is turned on cluster-wide automatically if the peer group being created results in multiple paths being available for an existing or future VIP interface.

**Required properties**

- `name` - Name of the peer-group to create.
- `ipspace.name` or `ipspace.uuid` or `local.port`
  - Required with local.interface.name or local.ip.address to identify a local interface

    - Optional when local.interface.uuid is specified
- `local.interface.uuid` or `local.interface.name` or `local.ip.address`
    - Required when specifying an existing local interface.
- `local.ip` and `local.port`
    - Required to create a new local interface.
- `peer.address` - IP address of the peer router

#### Default property values

If not specified in POST, the following default property values are assigned:

- `is_next_hop` - *false*
- `peer.md5_enabled` - *false*

#### Related ONTAP commands

- `network bgp peer-group create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Either the UUID or name is supplied on input. |
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |

| Name | Type | Description |
|---|---|---|
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

**Example request**

```
{
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local": {
    "interface": {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "ip": {
      "address": "10.10.10.7",
      "netmask": "24"
    },
    "port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7",
    "md5_enabled": 1,
    "md5_secret": "SECRET_WORD"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1376963 | Duplicate IP address is specified. |
| 1966133 | Since masking an address with a netmask represents an entire IP subnet, the masked and unmasked IP addresses cannot be the same. |
| 1966267 | IPv6 addresses must have a prefix length of 64. |
| 1966269 | IPv4 addresses must have a netmask length between 1 and 32. |
| 1967082 | IPspace name and UUID must match if both are given. |
| 1967155 | The specified local.port.name does not match the location.port.name for the specified local.interface. |
| 1967156 | The specified local.port.node.name does not match the location.port.node.name for the specified local.interface. |
| 1967157 | The specified local.port.uuid does not match the location.port.uuid for the specified local.interface. |
| 1967158 | The specified local.interface.name does not exist in the associated IPspace. local.ip.address and local.ip.netmask are required to create a new LIF. |
| 1967159 | local.interface does not support management-bgp service. |
| 1967160 | The specified local.interface.name does not match the specified interface name of local.interface.uuid. |
| 1967161 | The specified local.interface.uuid does not exist in the specified IPspace. |
| 1967162 | Either local.interface or local.ip and local.port are required to specify a local LIF. |
| 1967163 | The specified local.port.name does not match the specified port name of local.port.uuid. |

| Error Code | Description |
| --- | --- |
| 1967164 | The specified local.port.node.name does not match the specified node name of local.port.uuid. |
| 1967165 | The specified local.port does not exist. |
| 1967166 | ipspace.uuid or ipspace.name must be provided with local.interface.name together to identify a LIF. |
| 1967167 | Internal error. Failed to update BGP configuration for node. Retry the command, if necessary. |
| 1967168 | Internal error. Failed to create a VIP port for IPspace on node. Retry the command, if necessary. |
| 1967169 | Internal error. BGP configuration changed during the operation. Retry the command, if necessary. |
| 1967170 | Internal error. VIP port configuration changed during the operation. Retry the command, if necessary. |
| 1967171 | Internal error. Fail to access or update BGP peer group. Retry the command, if necessary. |
| 1967172 | Peer group could not be updated because IPspace does not exist. Retry the command, if necessary. |
| 1967173 | The specified local.ip.address does not match the address for the specified local.interface. |
| 1967174 | The specified local.ip.netmask does not match the netmask for the specified local.interface. |
| 1967176 | The specified local.interface.name does not exist in the associated IPspace. local.port.name, local.port.node.name, or local.port.uuid is required to create a new LIF. |
| 1967177 | Internal error. Failed to access the local interface. Retry the command, if necessary. |
| 1967178 | The IPv6 address specified with local.ip.address is not supported because it is link-local, multicast, v4-compatible, v4-mapped, loopback or "::". |
| 1967179 | The IPv4 address specified with local.ip.address is not supported because it is multicast, loopback or 0.0.0.0. |
| 1967187 | Configuring 4 bytes peer.asn requires an effective cluster version of 9.9.1 or later. |
| 1967188 | Configuring peer address as a next hop requires an effective cluster version of 9.9.1 or later. |
| 1967189 | The parameter peer.asn can’t be zero. |
| 1967192 | Configuring peer.md5_enabled requires an effective cluster version of 9.16.1 or later. |
| 1967193 | Configuring peer.md5_enabled requires the peer.md5_secret parameter. |

| Error Code | Description |
| --- | --- |
| 1967194 | Configuring peer.md5_secret requires that peer.md5_enabled is set to true. |
| 53281985 | Internal error. Failed to update BGP peer group because BGP LIF moved during the operation. Wait a few minutes and try the command again. |
| 53282006 | BGP peer group could not be updated to use a peer address because the value provided is not a valid peer address. If necessary, try the command again with a routable host address. |
| 53282007 | BGP peer group could not be updated to use a peer address because the address represents a different address family to the address of the associated BGP LIF. If necessary, try the command again with a matching address family. |
| 53282018 | Failed to create BGP peer group because an existing peer group has already established a BGP session between LIF and peer address. If necessary, try the command again with a different BGP LIF or a different peer address. |
| 53282021 | IPsec must be enabled to use TCP MD5 in BGP sessions with a BGP peer-group. |
| 53282024 | peer.md5_secret should be 1-80 bytes. |
| 53282025 | peer.md5_secret should have an even length. |
| 53282026 | The hex format of peer.md5_secret should have only hex string starting with '0x'. |
| 53282027 | peer.md5_secret should be 1-79 characters. |
| 53282028 | peer.md5_secret supports only printable ASCII characters such as numbers, alphabets, or special characters. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip

IP information

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ip

IP information to create a new interface.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

local

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |
| ip | ip | IP information to create a new interface. |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

peer

Information describing the router to peer with

| Name | Type | Description |
| --- | --- | --- |
| address | string | Peer router address |
| asn | integer | Autonomous system number of peer |
| is_next_hop | boolean | Use peer address as next hop. |
| md5_enabled | boolean | Enable or disable TCP MD5 signatures per RFC 2385. |
| md5_secret | string | The shared TCP MD5 secret key. This can either be given as a password or hexadecimal key. |

bgp_peer_group

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Either the UUID or name is supplied on input. |
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a BGP peer group for VIP

DELETE `/network/ip/bgp/peer-groups/{uuid}`

**Introduced In:** 9.7

Deletes a BGP peer group for VIP.

#### Related ONTAP commands

- `network bgp peer-group delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the peer group |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 53282019 | Internal error. Failed to remove BGP peer group on node. Wait a few minutes and try the command again. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve details of a BGP peer group for VIP

GET `/network/ip/bgp/peer-groups/{uuid}`

**Introduced In:** 9.7

Retrieves details of a BGP peer group for VIP.

#### Related ONTAP commands

- `network bgp peer-group show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the peer group |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Either the UUID or name is supplied on input. |
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

### Example response

```
{
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local": {
    "interface": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7",
    "md5_enabled": 1,
    "md5_secret": "SECRET_WORD"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

ip

IP information to create a new interface.

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

local

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | A network interface. Either UUID or name may be supplied on input. |

| Name | Type | Description |
| --- | --- | --- |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

peer

Information describing the router to peer with

| Name | Type | Description |
| --- | --- | --- |
| address | string | Peer router address |
| asn | integer | Autonomous system number of peer |
| is_next_hop | boolean | Use peer address as next hop. |
| md5_enabled | boolean | Enable or disable TCP MD5 signatures per RFC 2385. |
| md5_secret | string | The shared TCP MD5 secret key. This can either be given as a password or hexadecimal key. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a BGP peer group for VIP

PATCH `/network/ip/bgp/peer-groups/{uuid}`

**Introduced In:** 9.7

Updates a BGP peer group for VIP.

#### Related ONTAP commands

- `network bgp peer-group modify`
- `network bgp peer-group rename`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the peer group |

#### Request Body

| Name | Type | Description |
|---|---|---|
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

#### Example request

```
{
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7",
    "md5_enabled": 1,
    "md5_secret": "SECRET_WORD"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1967171 | Internal error. Fail to access or update BGP peer group. Retry the command, if necessary. |
| 1967188 | Configuring peer address as a next hop requires an effective cluster version of 9.9.1 or later. |
| 1967192 | Configuring peer.md5_enabled requires an effective cluster version of 9.16.1 or later. |
| 1967193 | Configuring peer.md5_enabled requires the peer.md5_secret parameter. |
| 1967194 | Configuring peer.md5_secret requires that peer.md5_enabled is set to true. |
| 53281998 | Failed to rename the BGP peer group because that name is already assigned to a different peer group in the IPspace. |
| 53282006 | BGP peer group could not be updated to use a peer address because the value provided is not a valid peer address. If necessary, try the command again with a routable host address. |

| Error Code | Description |
| --- | --- |
| 53282007 | BGP peer group could not be updated to use a peer address because the address represents a different address family to the address of the associated BGP LIF. If necessary, try the command again with a matching address family. |
| 53282018 | Failed to create BGP peer group because an existing peer group has already established a BGP session between LIF and peer address. If necessary, try the command again with a different BGP LIF or a different peer address. |
| 53282020 | To utilize the peer as the next hop, the peer IP address must be within the same subnet as the BGP LIF. |
| 53282021 | IPsec must be enabled to use TCP MD5 in BGP sessions with a BGP peer-group. |
| 53282024 | peer.md5_secret should be 1-80 bytes. |
| 53282025 | peer.md5_secret should have an even length. |
| 53282026 | The hex format of peer.md5_secret should only contain a hex string starting with '0x'. The specified peer.md5-secret contains non-hex characters. |
| 53282027 | peer.md5_secret should be 1-79 characters. |
| 53282028 | peer.md5_secret supports only printable ASCII characters such as numbers, alphabets, or special characters. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Either the UUID or name is supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip

IP information

interface

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ip

IP information to create a new interface.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

local

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

peer

Information describing the router to peer with

| Name | Type | Description |
| --- | --- | --- |
| address | string | Peer router address |
| is_next_hop | boolean | Use peer address as next hop. |
| md5_enabled | boolean | Enable or disable TCP MD5 signatures per RFC 2385. |
| md5_secret | string | The shared TCP MD5 secret key. This can either be given as a password or hexadecimal key. |

bgp_peer_group

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

| Name | Type | Description |
|---|---|---|
| local | local | Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address. |
| name | string | Name of the peer group |
| peer | peer | Information describing the router to peer with |
| state | string | State of the peer group |
| uuid | string | UUID of the peer group |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network IP interfaces

## Manage network IP interfaces

### Overview

The following operations are supported:

- Creation: POST network/ip/interfaces
- Collection Get: GET network/ip/interfaces
- Instance Get: GET network/ip/interfaces/{uuid}
- Instance Patch: PATCH network/ip/interfaces/{uuid}
- Instance Delete: DELETE network/ip/interfaces/{uuid}

### Retrieving network interface information

The IP interfaces GET API retrieves and displays relevant information pertaining to the interfaces configured in the cluster. The response can contain a list of multiple interfaces or a specific interface. The fields returned in the response vary for different interfaces and configurations.

### Examples

#### Retrieving all interfaces in the cluster

The following example shows the list of all interfaces configured in a cluster.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/interfaces" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "14531286-59fc-11e8-ba55-005056b4340f",
    "name": "user-cluster-01_mgmt1",
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/14531286-59fc-11e8-ba55-
005056b4340f"
      }
    }
  },
  {
    "uuid": "145318ba-59fc-11e8-ba55-005056b4340f",
```

```
    "name": "user-cluster-01_clus2",
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/145318ba-59fc-11e8-ba55-
005056b4340f"
      }
    }
  },
  {
    "uuid": "14531e45-59fc-11e8-ba55-005056b4340f",
    "name": "user-cluster-01_clus1",
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/14531e45-59fc-11e8-ba55-
005056b4340f"
      }
    }
  },
  {
    "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
    "name": "cluster_mgmt",
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
005056b4340f"
      }
    }
  },
  {
    "uuid": "c670707c-5a11-11e8-8fcb-005056b4340f",
    "name": "lif1",
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/c670707c-5a11-11e8-8fcb-
005056b4340f"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/network/ip/interfaces"
  }
}
}
```

### Retrieving a specific Cluster-scoped interface

The following example shows the response when a specific Cluster-scoped interface is requested. The system returns an error when there is no interface with the requested UUID. SVM information is not returned for Cluster-scoped interfaces.

```
# The API:
/api/network/ip/interfaces/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/interfaces/245979de-59fc-
11e8-ba55-005056b4340f" -H "accept: application/hal+json"

# The response:
{
"uuid": "245979de-59fc-11e8-ba55-005056b4340f",
"name": "cluster_mgmt",
"ip": {
  "address": "10.63.41.6",
  "netmask": "18",
  "family": "ipv4",
},
"enabled": true,
"state": "up",
"scope": "cluster",
"ipspace": {
  "uuid": "114ecfb5-59fc-11e8-ba55-005056b4340f",
  "name": "Default",
  "_links": {
    "self": {
            "href": "/api/network/ipspaces/114ecfb5-59fc-11e8-ba55-
005056b4340f"
    }
  }
},
"services": [
  "management_core",
  "management_autosupport",
  "management_access"
],
"location": {
  "is_home": true,
  "auto_revert": false,
  "failover": "broadcast_domain_only",
  "node": {
```

```
    "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
    "name": "user-cluster-01-a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-005056acfcbb"
      }
    }
  },
  "port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",
    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-
005056acfcbb"
      }
    }
  },
  "home_node": {
    "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
    "name": "user-cluster-01-a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-005056acfcbb"
      }
    }
  },
  "home_port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",
    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-
005056acfcbb"
      }
    }
  }
},
"service_policy": {
  "uuid": "9e0f4151-141b-11e9-851e-005056ac1ce0",
```

```
  "name": "default-management"
},
"vip": false,
"_links": {
  "self": {
    "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
005056b4340f"
  }
}
}
```

#### Retrieving a specific SVM-scoped interface using a filter

The following example shows the response when a specific SVM-scoped interface is requested. The SVM object is only included for SVM-scoped interfaces.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ip/interfaces?name=lif1&fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "c670707c-5a11-11e8-8fcb-005056b4340f",
    "name": "lif1",
    "ip": {
      "address": "10.10.10.11",
      "netmask": "24",
      "family": "ipv4",
    },
    "enabled": true,
    "state": "up",
    "scope": "svm",
    "ipspace": {
      "uuid": "114ecfb5-59fc-11e8-ba55-005056b4340f",
      "name": "Default",
      "_links": {
        "self": {
```

```
          "href": "/api/network/ipspaces/114ecfb5-59fc-11e8-ba55-
005056b4340f"
        }
      }
    },
    "svm": {
      "uuid": "c2134665-5a11-11e8-8fcb-005056b4340f",
      "name": "user_vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/c2134665-5a11-11e8-8fcb-005056b4340f"
        }
      }
    },
    "services": [
      "data_core",
      "data_nfs",
      "data_cifs",
      "data_flexcache"
    ],
    "location": {
      "is_home": true,
      "auto_revert": false,
      "failover": "broadcast_domain_only",
      "node": {
        "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
        "name": "user-cluster-01-a",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-
005056acfcbb"
          }
        }
      },
      "port": {
        "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
        "name": "e0d",
        "node": {
          "name": "user-cluster-01-a"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-
005056acfcbb"
          }
        }
```

```
      },
      "home_node": {
        "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
        "name": "user-cluster-01-a",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-
005056acfcbb"
          }
        }
      },
      "home_port": {
        "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
        "name": "e0d",
        "node": {
          "name": "user-cluster-01-a"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-
005056acfcbb"
          }
        }
      }
    },
    "service_policy": {
      "uuid": "9e53525f-141b-11e9-851e-005056ac1ce0",
      "name": "default-data-files"
    },
    "vip": false,
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/c670707c-5a11-11e8-8fcb-
005056b4340f"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/interfaces?name=lif1&fields=*"
  }
}
}
```

**Retrieving specific fields and limiting the output using filters**

The following example shows the response when a filter is applied (location.home_port.name=e0a) and only certain fields are requested. Filtered fields are in the output in addition to the default fields and requested fields.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ip/interfaces?location.home_port.name=e0a&fields=location.
home_node.name,service_policy.name,ip.address,enabled" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "1d1c9dc8-4f17-11e9-9553-005056ac918a",
    "name": "user-cluster-01-a_clus1",
    "ip": {
      "address": "192.168.170.24"
    },
    "enabled": true,
    "location": {
      "home_node": {
        "name": "user-cluster-01-a"
      },
      "home_port": {
        "name": "e0a"
      }
    },
    "service_policy": {
      "name": "default-cluster"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/1d1c9dc8-4f17-11e9-9553-
005056ac918a"
      }
    }
  },
  {
    "uuid": "d07782c1-4f16-11e9-86e7-005056ace7ee",
```

```
    "name": "user-cluster-01-b_clus1",
    "ip": {
      "address": "192.168.170.22"
    },
    "enabled": true,
    "location": {
      "home_node": {
        "name": "user-cluster-01-b"
      },
      "home_port": {
        "name": "e0a"
      }
    },
    "service_policy": {
      "name": "default-cluster"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/d07782c1-4f16-11e9-86e7-
005056ace7ee"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/network/ip/interfaces?location.home_port.name=e0a&fields=location.ho
me_node.name,service_policy.name,ip.address,enabled"
  }
}
}
```


#### Creating IP interfaces

You can use the IP interfaces POST API to create IP interfaces as shown in the following examples.

#### Examples

### Creating a Cluster-scoped IP interface using names

The following example shows the record returned after the creation of an IP interface on "e0d".

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "cluster_mgmt", "ip": { "address":
"10.63.41.6", "netmask": "18" }, "enabled": true, "scope": "cluster",
"ipspace": { "name": "Default" }, "location": { "auto_revert": false,
"failover": "broadcast_domain_only", "home_port": { "name": "e0d", "node":
{ "name": "user-cluster-01-a" } } }, "service_policy": { "name": "default-
management" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
    "name": "cluster_mgmt",
    "ip": {
      "address": "10.63.41.6",
      "netmask": "18"
    },
    "enabled": true,
    "scope": "cluster",
    "ipspace": {
      "name": "Default"
    },
    "location": {
      "auto_revert": false,
      "failover": "broadcast_domain_only",
      "home_port": {
        "name": "e0d",
        "node": {
          "name": "user-cluster-01-a"
        }
      },
    },
    "service_policy": {
      "name": "default-management"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
```

```
005056b4340f"
      }
    }
  }
]
}
```

#### Creating a SVM-scoped IP interface using a mix of parameter types

The following example shows the record returned after the creation of a IP interface by specifying a broadcast domain as the location.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "Data1", "ip": { "address":
"10.234.101.116", "netmask": "255.255.240.0" }, "enabled": true, "scope":
"svm", "svm": { "uuid": "137f3618-1e89-11e9-803e-005056a7646a" },
"location": { "auto_revert": true, "broadcast_domain": { "name": "Default"
} }, "service_policy": { "name": "default-data-files" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "80d271c9-1f43-11e9-803e-005056a7646a",
    "name": "Data1",
    "ip": {
      "address": "10.234.101.116",
      "netmask": "20"
    },
    "enabled": true,
    "scope": "svm",
    "svm": {
      "uuid": "137f3618-1e89-11e9-803e-005056a7646a",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/137f3618-1e89-11e9-803e-005056a7646a"
```

```
        }
      }
    },
    "location": {
      "auto_revert": true
    },
    "service_policy": {
      "name": "default-data-files"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-
005056a7646a"
      }
    }
  }
]
}
```


#### Creating a Cluster-scoped IP interface without specifying the scope parameter

The following example shows the record returned after creating an IP interface on "e0d" without specifying the scope parameter. The scope is "cluster" if an "svm" is not specified.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "cluster_mgmt", "ip": { "address":
"10.63.41.6", "netmask": "18" }, "enabled": true, "ipspace": { "name":
"Default" }, "location": { "auto_revert": false, "home_port": { "name":
"e0d", "node": { "name": "user-cluster-01-a" } } }, "service_policy": {
"name": "default-management" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
    "name": "cluster_mgmt",
```

```
    "ip": {
      "address": "10.63.41.6",
      "netmask": "18"
    },
    "enabled": true,
    "scope": "cluster",
    "ipspace": {
      "name": "Default"
    },
    "location": {
      "auto_revert": false,
      "home_port": {
        "name": "e0d",
        "node": {
          "name": "user-cluster-01-a"
        }
      }
    },
    "service_policy": {
      "name": "default-management"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
005056b4340f"
      }
    }
  }
]
}
```


#### Creating an SVM-scoped IP interface without specifying the scope parameter

The following example shows the record returned after creating an IP interface on "e0d" without specifying the scope parameter. The scope is "svm" if the "svm" field is specified.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "Data1", "ip": { "address":
```

```
"10.234.101.116", "netmask": "255.255.240.0" }, "enabled": true, "svm": {
"uuid": "137f3618-1e89-11e9-803e-005056a7646a" }, "location": {
"auto_revert": true, "broadcast_domain": { "name": "Default" } },
"service_policy": { "name": "default-data-files" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "80d271c9-1f43-11e9-803e-005056a7646a",
    "name": "Data1",
    "ip": {
      "address": "10.234.101.116",
      "netmask": "20"
    },
    "enabled": true,
    "scope": "svm",
    "svm": {
      "uuid": "137f3618-1e89-11e9-803e-005056a7646a",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svms/137f3618-1e89-11e9-803e-005056a7646a"
        }
      }
    },
    "location": {
      "auto_revert": true
    },
    "service_policy": {
      "name": "default-data-files"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-
005056a7646a"
      }
    }
  }
]
}
```

#### Creating an SVM-scoped IP interface using a subnet

The following example shows the record returned after the creation of a IP interface by allocating an IP address from a subnet.

```
# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "Data1", "subnet": { "name":
"Subnet10" }, "enabled": true, "scope": "svm", "svm": { "uuid": "137f3618-
1e89-11e9-803e-005056a7646a" }, "location": { "auto_revert": true,
"broadcast_domain": { "name": "Default" } }, "service_policy": { "name":
"default-data-files" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "80d271c9-1f43-11e9-803e-005056a7646a",
    "name": "Data1",
    "enabled": true,
    "scope": "svm",
    "svm": {
      "uuid": "137f3618-1e89-11e9-803e-005056a7646a",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/137f3618-1e89-11e9-803e-005056a7646a"
        }
      }
    },
    "location": {
      "auto_revert": true
    },
    "service_policy": {
      "name": "default-data-files"
    },
    "subnet": {
      "name": "testSubnet"
    },
    "_links": {
```

```
      "self": {
        "href": "/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-
005056a7646a"
      }
    }
  }
]
}
```

#### Updating IP interfaces

You can use the IP interfaces PATCH API to update the attributes of an IP interface.

#### Examples

### Updating the auto revert flag of an IP interface

The following example shows how the PATCH request changes the auto revert flag to 'false'.

```
# The API:
/api/network/ip/interfaces/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-
11e9-803e-005056a7646a" -H "accept: application/hal+json" -d '{
"location": { "auto_revert": "false" } }'
{
}
```

### Updating the service policy of an IP interface

The following example shows how the PATCH request changes the service policy to 'default-management'.

```
# The API:
/api/network/ip/interfaces/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-
11e9-803e-005056a7646a" -H "accept: application/hal+json" -d '{
"service_policy": { "name": "default-management" }}'
{
}
```

#### Deleting IP interfaces

You can use the IP interfaces DELETE API to delete an IP interface in the cluster.

#### Example

### Deleting an IP Interface

The following DELETE request deletes a network IP interface.

```
# The API:
/api/network/ip/interfaces/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-
11e9-803e-005056a7646a"
{
}
```

### Retrieve all IP interface details

GET `/network/ip/interfaces`

**Introduced In:** 9.6

Retrieves the details of all IP interfaces.

#### Related ONTAP Commands

- `network interface show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| vip | boolean | query | False | Filter by vip |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| probe_port | integer | query | False | Filter by probe_port<br>• Introduced in: 9.10 |
| scope | string | query | False | Filter by scope |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| rdma_protocols | string | query | False | Filter by rdma_protocols<br>• Introduced in: 9.10 |
| service_policy.uuid | string | query | False | Filter by service_policy.uuid |
| service_policy.name | string | query | False | Filter by service_policy.name |
| location.auto_revert | boolean | query | False | Filter by location.auto_revert |
| location.is_home | boolean | query | False | Filter by location.is_home |
| location.failover | string | query | False | Filter by location.failover |
| location.port.node.name | string | query | False | Filter by location.port.node.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| location.port.uuid | string | query | False | Filter by location.port.uuid |
| location.port.name | string | query | False | Filter by location.port.name |
| location.home_node.name | string | query | False | Filter by location.home_node.name |
| location.home_node.uuid | string | query | False | Filter by location.home_node.uuid |
| location.node.name | string | query | False | Filter by location.node.name |
| location.node.uuid | string | query | False | Filter by location.node.uuid |
| location.home_port.node.name | string | query | False | Filter by location.home_port.node.name |
| location.home_port.uuid | string | query | False | Filter by location.home_port.uuid |
| location.home_port.name | string | query | False | Filter by location.home_port.name |
| services | string | query | False | Filter by services |
| enabled | boolean | query | False | Filter by enabled |
| name | string | query | False | Filter by name |
| subnet.name | string | query | False | Filter by subnet.name<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| subnet.uuid | string | query | False | Filter by subnet.uuid<br>• Introduced in: 9.11 |
| state | string | query | False | Filter by state |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br>• Introduced in: 9.8 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.8 |
| ip.family | string | query | False | Filter by ip.family |
| ip.address | string | query | False | Filter by ip.address |
| ip.netmask | string | query | False | Filter by ip.netmask |
| uuid | string | query | False | Filter by uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| ddns_enabled | boolean | query | False | Filter by ddns_enabled<br>• Introduced in: 9.9 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.8 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br>• Introduced in: 9.8 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.8 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.8 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br>• Introduced in: 9.8 |
| dns_zone | string | query | False | Filter by dns_zone<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.svm.uuid | string | query | False | The UUID of an existing SVM for which to recommend IP interfaces.<br><br>• Introduced in: 9.16 |
| recommend.svm.name | string | query | False | The name of an existing SVM for which to recommend IP interfaces.<br><br>• Introduced in: 9.16 |
| recommend.ipspace.uuid | string | query | False | The UUID of an existing IPspace for which to recommend IP interfaces.<br><br>• Introduced in: 9.16 |
| recommend.ipspace.name | string | query | False | The name of an existing IPspace for which to recommend IP interfaces.<br><br>• Introduced in: 9.16 |
| recommend.service_policy.uuid | string | query | False | The UUID of an IP service policy for which to recommend IP interfaces.<br><br>• Introduced in: 9.16 |
| recommend.service_policy.name | string | query | False | The name of an IP service policy for which to recommend IP interfaces.<br><br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| recommend.ip.family | string | query | False | The name of an IP Address Family for which to recommend IP interfaces. Unless ipv6 is specified, the default value is ipv4.<br>• Introduced in: 9.16<br>• enum: ["ipv4", "ipv6"] |
| recommend.location.broadcast_domain.uuid | string | query | False | The UUID of an existing broadcast domain for which to recommend IP interfaces.<br>• Introduced in: 9.16 |
| recommend.location.broadcast_domain.name | string | query | False | The name of an existing broadcast domain for which to recommend IP interfaces.<br>• Introduced in: 9.16 |
| recommend.subnet.uuid | string | query | False | The UUID of an existing subnet for which to recommend IP interfaces.<br>• Introduced in: 9.16 |
| recommend.subnet.name | string | query | False | The name of an existing subnet for which to recommend IP interfaces.<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| recommend.interfaces_per_node | integer | query | False | The number of IP interfaces to recommend per node.<br><br>• Introduced in: 9.16 |
| recommend.location.port.rdma_protocol | string | query | False | The RDMA protocol supported on locations for which to recommend IP interfaces.<br><br>• Introduced in: 9.17 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br><br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| recommend | recommend | Response properties specific to the LIF recommendation functionality. |
| records | array[ip_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "recommend": {
    "messages": [
      {}
    ]
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "dns_zone": "storage.company.com",
      "fail_if_subnet_conflicts": true,
      "ip": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "location": {
        "failover": "string",
        "home_node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "home_port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
```

```
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "dataLif1",
      "probe_port": 64001,
      "rdma_protocols": [
        "roce"
      ],
      "scope": "string",
      "service_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default-intercluster",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "services": [
        "data_nfs"
      ],
      "state": "string",
      "statistics": {
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "subnet": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "subnet1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 53282769 | Not all nodes support enabling the LIF recommendations feature. |
| 53282770 | No matching subnets are configured in IPspace. |
| 53282772 | Not enough available IPs in subnet. |
| 53282773 | Not enough locations available on node for subnet. |
| 53282774 | The specified SVM UUID is invalid. |
| 53282775 | The specified SVM name is invalid. |
| 53282776 | The specified SVM name and UUID do not match. |
| 53282777 | The specified SVM does not exist in the IPspace. |
| 53282778 | The specified SVM does not exist in the IPspace. |
| 53282779 | The specified IPspace UUID is invalid. |
| 53282780 | The specified IPspace name is invalid. |
| 53282781 | The specified IPspace UUID and name do not match. |
| 53282782 | The specified service policy UUID does not exist. |
| 53282783 | THe specified service policy UUID and name do not match. |
| 53282784 | The specified service policy does not exist in the SVM. |

| Error Code | Description |
|---|---|
| 53282785 | The specified service policy does not exist in the IPspace. |
| 53282786 | The specified broadcast domain does not exist. |
| 53282787 | The specified broadcast domain UUID and name do not match. |
| 53282788 | The specified broadcast domain UUID does not exist in the IPspace. |
| 53282789 | The specified broadcast domain name does not exist in the IPspace. |
| 53282790 | The specified subnet UUID does not exist. |
| 53282791 | The specified subnet UUID and name do not match. |
| 53282792 | The specified subnet does not exist in the IPspace. |
| 53282793 | The specified subnet does not match the specified broadcast domain. |
| 53282794 | The specified subnet name does not exist in the IPspace. |
| 53282795 | The specified subnet name does not exist in the broadcast domain. |
| 53282796 | The specified SVM has a type that is not supported. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

ip_interface_recommend_message

recommend

Response properties specific to the LIF recommendation functionality.

| Name | Type | Description |
|---|---|---|
| messages | array[ip_interface_recommend_message] | Messages describing the results of a LIF recommendation request. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

home_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

location

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port.

| Name | Type | Description |
|---|---|---|
| auto_revert | boolean | |
| failover | string | Policy that defines where an interface is permitted to move on failover. The policy named “default” implements the recommended best practice for NAS LIFs on the current platform and cluster, and was known as “system_defined” in the CLI. |
| home_node | home_node | |
| home_port | home_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| is_home | boolean | |
| node | node | |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

service_policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

throughput_raw

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_interface

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |
| ip | ip_info | IP information |
| ipspace | ipspace | Either the UUID or name must be supplied on POST for cluster-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| metric | metric | The most recent sample of I/O metrics for the interface. |
| name | string | Interface name |
| probe_port | integer | Probe port for Cloud load balancer |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| statistics | statistics | The real time I/O statistics for the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID that uniquely identifies the interface. |
| vip | boolean | True for a VIP interface, whose location is announced via BGP. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new cluster-scoped or SVM-scoped interface**

POST `/network/ip/interfaces`

**Introduced In:** 9.6

Creates a new Cluster-scoped or SVM-scoped interface.

**Required properties**

- `name` - Name of the interface to create.
- `ip` or `subnet`
  - `ip.address` - IP address for the interface.
  - `ip.netmask` - IP subnet of the interface.
  - `subnet.uuid` or `subnet.name`
- `ipspace.name` or `ipspace.uuid`

    - Required for Cluster-scoped interfaces.
    - Optional for SVM-scoped interfaces.
- `svm.name` or `svm.uuid`
    - Required for an SVM-scoped interface.
    - Invalid for a Cluster-scoped interface.
- If a LIF in the subnet of the specified IP address does not already exist, a location.home_port, a location.home_node, or a location.broadcast_domain needs to be provided.

#### Recommended property values

- `service_policy`
    - `for SVM scoped interfaces`
        - *default-data-files* for interfaces carrying file-oriented NAS data traffic
        - *default-data-blocks* for interfaces carrying block-oriented SAN data traffic
        - *default-data-iscsi* for interfaces carrying iSCSI data traffic
        - *default-management* for interfaces carrying SVM management requests
    - `for Cluster scoped interfaces`
        - *default-intercluster* for interfaces carrying cluster peering traffic
        - *default-management* for interfaces carrying system management requests
        - *default-route-announce* for interfaces carrying BGP peer connections

#### Default property values

If not specified in POST, the following default property values are assigned:

- `scope`
    - *svm* if svm parameter is specified.
    - *cluster* if svm parameter is not specified
- `enabled` - *true*
- `location.auto_revert` - *true*
- `service_policy`
    - ***Unified ONTAP**: *default-data-files* if scope is `svm`
    - ***ASA r2**: *default-data-blocks* if scope is `svm`
        - *default-management* if scope is `cluster` and IPspace is not `Cluster`
        - *default-cluster* if scope is `cluster` and IPspace is `Cluster`
- `failover` - Selects the least restrictive failover policy supported by all the services in the service policy.
- `ddns_enabled`
    - *true* if the interface supports *data_nfs* or *data_cifs* services
    - *false* otherwise

- `fail_if_subnet_conflicts` - *true*

#### Related ONTAP commands

- `network interface create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |
| fail_if_subnet_conflicts | boolean | This command fails if the specified IP address falls within the address range of a named subnet. Set this value to false to use the specified IP address and to assign the subnet owning that address to the interface. |
| ip | ip_info | IP information |
| ipspace | ipspace | Either the UUID or name must be supplied on POST for cluster-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| name | string | Interface name |
| probe_port | integer | Probe port for Cloud load balancer |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST. |
| uuid | string | The UUID that uniquely identifies the interface. |
| vip | boolean | True for a VIP interface, whose location is announced via BGP. |

**Example request**


```
{
  "dns_zone": "storage.company.com",
  "ip": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "broadcast_domain": {
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "failover": "string",
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "is_home": true
  },
  "name": "dataLif1",
  "probe_port": 64001,
  "rdma_protocols": [
    "roce"
  ],
  "scope": "string",
  "service_policy": {
    "name": "default-intercluster",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "services": [
    "data_nfs"
  ],
  "state": "string",
  "subnet": {
```

```
      "name": "subnet1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| recommend | recommend | Response properties specific to the LIF recommendation functionality. |
| records | array[ip_interface] | |

**Example response**


```
{
  "num_records": 1,
  "recommend": {
    "messages": [
      {}
    ]
  },
  "records": [
    {
      "dns_zone": "storage.company.com",
      "ip": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "location": {
        "broadcast_domain": {
          "name": "bd1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "failover": "string",
        "home_node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "home_port": {
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "is_home": true
      },
      "name": "dataLif1",
      "probe_port": 64001,
      "rdma_protocols": [
        "roce"
      ],
      "scope": "string",
      "service_policy": {
```

```
        "name": "default-intercluster",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "services": [
        "data_nfs"
      ],
      "state": "string",
      "subnet": {
        "name": "subnet1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1376656 | Cluster interfaces must be in the same subnet. Verify the address and netmask are set to the correct values. |
| 1376663 | All LIFs from a single DNS zone must be in the same SVM. |
| 1376663 | Cannot add interface to DNS zone because all interfaces from a single DNS zone must be in the same SVM. |
| 1376963 | Duplicate IP address. |
| 1376976 | The specified port is not capable of hosting this LIF. |

| Error Code | Description |
| --- | --- |
| 1377583 | Failed to create the interface because the broadcast domain of the specified subnet is different from the specified broadcast domain. |
| 1377666 | Subnet does not have any addresses available. |
| 1966138 | The same IP address may not be used for both a mgmt interface and a gateway address. |
| 1966140 | An interface with the same name already exists. |
| 1966141 | Invalid DNS zone name. |
| 1966142 | Only data LIFs can be assigned a DNS zone. |
| 1966191 | The interface could not be created because interface identifier creation failed. |
| 1966217 | Invalid port on the node. |
| 1966267 | IPv6 addresses must have a prefix length between 1 and 127. |
| 1966269 | IPv4 addresses must have a prefix length between 1 and 32. |
| 1966270 | Operation not support on SAN LIFs. |
| 1966300 | The LIF name is too long. The maximum number of characters allowed for iSCSI and FC LIF names is 254. |
| 1966373 | Failed to create interface because the home-port is not in the IPspace associated with the SVM. |
| 1966454 | A port on the node is not a member of a broadcast domain. |
| 1966476 | DNS Update is supported only on data LIFs. |
| 1966477 | DNS Update is supported only on LIFs configured with the NFS or CIFS protocol. |
| 1966987 | The Vserver Broadcast-Domain Home-Node and Home-Port combination is not valid. |
| 1967081 | The specified SVM must exist in the specified IPspace. |
| 1967082 | The specified ipspace.name does not match the IPspace name of ipspace.uuid. |
| 1967102 | POST operation might have left configuration in an inconsistent state. Check the configuration. |
| 1967106 | The specified location.home_port.name does not match the specified port name of location.home_port.uuid. |
| 1967107 | The location.home_port.uuid specified is not valid. |

| Error Code | Description |
| --- | --- |
| 1967108 | The specified location.home_node.name does not match the node name of location.home_node.uuid. |
| 1967109 | The specified location.home_port.node.name does not match the node name of location.home_node.uuid. |
| 1967110 | The specified location.home_port.node.name does not match location.home_node.name. |
| 1967111 | Home node must be specified by at least one location.home_node, location.home_port, or location.broadcast_domain field. |
| 1967112 | The specified location.home_node.name does not match the node name of location.home_port.uuid. |
| 1967120 | The specified service_policy.name does not match the specified service policy name of service_policy.uuid. |
| 1967121 | Invalid service_policy.uuid specified. |
| 1967122 | The specified location.broadcast_domain.name does not match the specified broadcast domain name of location.broadcast_domain.uuid. |
| 1967123 | The specified IPspace does not match the IPspace name of location.broadcast_domain.uuid. |
| 1967124 | The location.broadcast_domain.uuid specified is not valid. |
| 1967127 | svm.uuid or svm.name must be provided if scope is "svm". |
| 1967128 | ipspace.uuid or ipspace.name must be provided if scope is "cluster". |
| 1967129 | The specified location.home_port.uuid is not valid. |
| 1967130 | The specified location.home_port.name is not valid. |
| 1967131 | The specified location.home_port.uuid and location.home_port.name are not valid. |
| 1967135 | The specified location.broadcast_domain.uuid is not valid. |
| 1967136 | The specified location.broadcast_domain.name (and ipspace name) is not valid. |
| 1967137 | The specified location.broadcast_domain.uuid and location.broadcast_domain.name (and IPspace name) are not valid. |
| 1967145 | The specified location.failover is not valid. |
| 1967146 | The specified svm.name is not valid. |
| 1967147 | The specified svm.uuid is not valid. |

| Error Code | Description |
| --- | --- |
| 1967153 | No suitable port exists on location.home_node to host the interface. |
| 1967154 | Interfaces cannot be created on ports that are down. If a broadcast domain is specified, ensure that it contains at least one port that is operationally up. |
| 1967381 | Post VIP interfaces requires an effective cluster version of 9.7 or later. |
| 1967382 | VIP interfaces only reside in SVM scope. |
| 1967383 | Neither location.home_port.uuid or location.home_port.name should be set with vip=true. |
| 1967384 | Failed to create VIP interface because the home node does not have active BGP sessions to support Virtual IP (VIP) traffic. |
| 1967385 | VIP interfaces with an IPv4 address must use ip.netmask=32. VIP interfaces with an IPv6 address must use ip.netmask=128. |
| 1967387 | The specified IP address is in use by a subnet in this IPspace. |
| 1967391 | Setting the DNS zone requires an effective cluster version of 9.9.1 or later. |
| 1967392 | Setting the DDNS enable parameter requires an effective cluster version of 9.9.1 or later. |
| 1967394 | Setting the probe port parameter requires an effective cluster version of 9.10.1 or later. |
| 1967396 | The specified subnet.name does not match the subnet name of subnet.uuid. |
| 1967397 | The specified subnet.uuid does not match any configured subnet."; |
| 1967398 | Address must be specified by either ip.address and ip.netmask, or at least one subnet field, not both."; |
| 1967401 | The specified subnet.name does not match any configured subnet. |
| 1967402 | Data SVM cannot be specified with cluster scope. |
| 2621519 | Invalid SVM name. The SVM name must begin with a letter or an underscore. The maximum supported length is 41 if the SVM type is "sync-source", otherwise it is 47. |
| 5373966 | An iSCSI interface cannot be created in an SVM configured for NVMe. |
| 8847378 | Cannot specify DNS zone when DNS updates are enabled for SVM |

| Error Code | Description |
| --- | --- |
| 53216540 | LIFs on the specified SVM do not have a default service policy. |
| 53281018 | Failover policy is not compatible with one or more services in service policy |
| 53281036 | Setting the probe port parameter is not allowed on this platform. |
| 53281065 | The service_policy does not exist in the SVM. |
| 53281073 | IP address and netmask cannot be used because this represents a subnet address rather than a specific host address. |
| 53281086 | LIF would exceed the maximum number of supported intercluster LIFs in IPspace. |
| 53281087 | Cannot configure SAN LIF on SVM. |
| 53281092 | Failed to create interface because the home-port does not support the specified protocol. |
| 53281104 | The specified address is in use by the Service Processor |
| 53281106 | Failed checking the cluster capabilities. |
| 53281114 | The specified parameter is only supported on data SVMs. |
| 53281468 | Failed to update external route tables for IP. |
| 53281680 | Invalid value for the specified field. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
|---|---|---|
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

home_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

location

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port.

| Name | Type | Description |
| --- | --- | --- |
| auto_revert | boolean | |

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. |
| failover | string | Policy that defines where an interface is permitted to move on failover. The policy named “default” implements the recommended best practice for NAS LIFs on the current platform and cluster, and was known as “system_defined” in the CLI. |
| home_node | home_node | |
| home_port | home_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

service_policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

throughput_raw

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_interface

| Name | Type | Description |
| --- | --- | --- |
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |
| fail_if_subnet_conflicts | boolean | This command fails if the specified IP address falls within the address range of a named subnet. Set this value to false to use the specified IP address and to assign the subnet owning that address to the interface. |
| ip | ip_info | IP information |

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Either the UUID or name must be supplied on POST for cluster-scoped objects. |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| name | string | Interface name |
| probe_port | integer | Probe port for Cloud load balancer |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST. |
| uuid | string | The UUID that uniquely identifies the interface. |

| Name | Type | Description |
|---|---|---|
| vip | boolean | True for a VIP interface, whose location is announced via BGP. |

ip_interface_recommend_message

recommend

Response properties specific to the LIF recommendation functionality.

| Name | Type | Description |
|---|---|---|
| messages | array[ip_interface_recommend_message] | Messages describing the results of a LIF recommendation request. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an IP interface

DELETE `/network/ip/interfaces/{uuid}`

**Introduced In:** 9.6

Deletes an IP interface.

#### Related ONTAP commands

- `network interface delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP interface UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966465 | LIF cannot be removed because it is required to maintain quorum on the node. |
| 5376953 | The interface is part of one or more portsets. Remove the interface from all portsets before deleting it. |
| 53281112 | Failed to delete the interface because it has an associated BGP peer group. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve details for an IP interface

GET `/network/ip/interfaces/{uuid}`

**Introduced In:** 9.6

Retrieves details for a specific IP interface.

#### Related ONTAP commands

- `network interface show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP interface UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |
| ip | ip_info | IP information |
| ipspace | ipspace | Either the UUID or name must be supplied on POST for cluster-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| metric | metric | The most recent sample of I/O metrics for the interface. |
| name | string | Interface name |
| probe_port | integer | Probe port for Cloud load balancer |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| statistics | statistics | The real time I/O statistics for the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST. |
| uuid | string | The UUID that uniquely identifies the interface. |

| Name | Type | Description |
| --- | --- | --- |
| vip | boolean | True for a VIP interface, whose location is announced via BGP. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "dns_zone": "storage.company.com",
  "fail_if_subnet_conflicts": true,
  "ip": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "failover": "string",
    "home_node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "e1b",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "metric": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "name": "dataLif1",
    "probe_port": 64001,
    "rdma_protocols": [
      "roce"
    ],
    "scope": "string",
    "service_policy": {
      "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default-intercluster",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "services": [
    "data_nfs"
  ],
  "state": "string",
  "statistics": {
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "subnet": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "subnet1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

home_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

location

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port.

| Name | Type | Description |
| --- | --- | --- |
| auto_revert | boolean | |
| failover | string | Policy that defines where an interface is permitted to move on failover. The policy named “default” implements the recommended best practice for NAS LIFs on the current platform and cluster, and was known as “system_defined” in the CLI. |
| home_node | home_node | |
| home_port | home_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| is_home | boolean | |
| node | node | |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

service_policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

throughput_raw

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope must be provided by the object this object is embedded in. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID that uniquely identifies the subnet. |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update an IP interface**

PATCH `/network/ip/interfaces/{uuid}`

**Introduced In:** 9.6

Updates an IP interface.

#### Related ONTAP commands

- `network interface migrate`
- `network interface modify`
- `network interface rename`
- `network interface revert`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP interface UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |
| fail_if_subnet_conflicts | boolean | This command fails if the specified IP address falls within the address range of a named subnet. Set this value to false to use the specified IP address and to assign the subnet owning that address to the interface. |
| ip | ip_info | IP information |

| Name | Type | Description |
| --- | --- | --- |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| name | string | Interface name |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

**Example request**


```
{
  "dns_zone": "storage.company.com",
  "ip": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "location": {
    "failover": "string",
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "dataLif1",
  "rdma_protocols": [
    "roce"
  ],
  "service_policy": {
    "name": "default-intercluster",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "services": [
    "data_nfs"
  ],
  "state": "string",
  "subnet": {
```

```
      "name": "subnet1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "vip": true
  }
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262196 | Field cannot be set in this operation. |
| 1376663 | Cannot add interface to DNS zone because all interfaces from a single DNS zone must be in the same SVM. |
| 1376963 | Duplicate IP address. |
| 1376976 | The specified port is not capable of hosting this LIF. |
| 1376986 | The interface could not migrate because no additional interfaces can be hosted on the specified node. |
| 1376997 | Interface failed to migrate because the node hosting the port is not healthy. |
| 1376998 | The specified location.node does not own any ports in the same broadcast domain as the home port of the interface. |
| 1376999 | Interface failed to migrate because port is in the down admin state. |
| 1377607 | The specified location.port is not in the same broadcast domain as the home port of the interface. |
| 1377666 | Subnet does not have any addresses available. |
| 1966081 | Failed to lookup the port on the node. |
| 1966133 | The netmask cannot represent the entire IP subnet. |
| 1966135 | Cluster configuration can only be changed from node where the cluster LIF resides. |

| Error Code | Description |
| --- | --- |
| 1966138 | The same IP address may not be used for both a mgmt interface and a gateway address. |
| 1966141 | Invalid DNS zone name. |
| 1966142 | Only data LIFs can be assigned a DNS zone. |
| 1966191 | The interface could not be created because interface identifier creation failed. |
| 1966197 | Migration of cluster interfaces must be done from the local node. |
| 1966238 | Cannot change the home-node or home-port of an active SAN or NVMe data interface. |
| 1966267 | IPv6 addresses must have a prefix length between 1 and 127. |
| 1966269 | IPv4 addresses must have a prefix length between 1 and 32. |
| 1966299 | Changing a cluster LIF's administrative status to down is not a supported operation. |
| 1966419 | Subnet not found in the IPspace of the SVM. |
| 1966419 | Specified subnet not found in the IPspace of the LIF. |
| 1966476 | DNS Update is supported only on data interfaces. |
| 1966477 | DNS Update is supported only on interfaces configured with the NFS or CIFS protocol. |
| 1966507 | A port on the node has been identified as potentially unhealthy. |
| 1966685 | There are no ports on the node that are in the broadcast domain associated with the IPspace. |
| 1967106 | The specified location.home_port.name does not match the specified port name of location.home_port.uuid. |
| 1967107 | The specified location.home_port.uuid is not valid. |
| 1967111 | A home node must be specified by at least one location.home_node, location.home_port, or location.broadcast_domain field. |
| 1967113 | The specified location.port.name does not match the port name of location.port.uuid. |
| 1967114 | The specified location.port.uuid is not valid. |
| 1967115 | The specified location.node.name does not match the node name of location.node.uuid. |
| 1967116 | The specified location.port.node.name does not match the node name of location.node.uuid. |

| Error Code | Description |
| --- | --- |
| 1967117 | The specified location.port.node.name does not match location.node.name. |
| 1967118 | A node must be specified by at least one location.node or location.port field. |
| 1967119 | The specified location.node.name does not match the node name of location.port.uuid. |
| 1967120 | The specified service_policy.name does not match the specified service policy name of service_policy.uuid. |
| 1967121 | The specified service_policy.uuid is not valid. |
| 1967125 | You cannot patch the "location.node" or "location.port" fields to migrate interfaces using the iSCSI data protocol. Instead perform the following PATCH operations on the interface: set the "enabled" field to "false"; change one or more "location.home_port" fields to migrate the interface; and then set the "enabled" field to "true". |
| 1967129 | The specified location.home_port.uuid is not valid. |
| 1967130 | The specified location.home_port.name is not valid. |
| 1967131 | The specified location.home_port.uuid and location.home_port.name are not valid. |
| 1967132 | The specified location.port.uuid is not valid. |
| 1967133 | The specified location.port.name is not valid. |
| 1967134 | The specified location.port.uuid and location.port.name are not valid. |
| 1967138 | Cannot patch port for a VIP interface. The specified parameter location.port.uuid is not valid. |
| 1967139 | Cannot patch port for a VIP interface. The specified parameter location.port.name is not valid. |
| 1967140 | Cannot patch port for a VIP interface. The specified parameters location.port.uuid and location.port.name are not valid. |
| 1967141 | Cannot patch home_port for a VIP interface. The specified parameter location.home_port.uuid is not valid. |
| 1967142 | Cannot patch home_port for a VIP interface. The specified parameter location.home_port.name is not valid. |
| 1967143 | Cannot patch home_port for a VIP interface. The specified parameters location.home_port.uuid and location.home_port.name are not valid. |
| 1967145 | The specified location.failover is not valid. |

| Error Code | Description |
| --- | --- |
| 1967153 | No suitable port exists on location.home_node to host the interface. |
| 1967380 | Cannot patch home_port for a VIP interface. The specified parameter location.home_port.node.name is not valid. Consider using location.home_node.name instead. |
| 1967385 | VIP interface address and netmask error. |
| 1967386 | Cannot patch port for a VIP interface. The specified parameter location.port.node.name is not valid. Consider using location.node.name instead. |
| 1967387 | The specified IP address is in use by a subnet in this IPspace. |
| 1967389 | Patching location.is_home to the value "false" is not supported. The value "true" would revert a network interface to its home port if the current value is "false". |
| 1967390 | Cannot patch a LIF revert as it requires an effective cluster version of 9.9.1 or later. |
| 1967391 | Patching the DNS zone requires an effective cluster version of 9.9.1 or later. |
| 1967392 | Patching the DDNS enable parameter requires an effective cluster version of 9.9.1 or later. |
| 1967396 | The specified subnet.name does not match the subnet name of subnet.uuid. |
| 1967397 | The specified subnet.uuid does not match any configured subnet."; |
| 1967398 | Address must be specified by either ip.address and ip.netmask, or at least one subnet field, not both."; |
| 2621574 | This operation is not permitted on an SVM that is configured as the destination of a MetroCluster SVM relationship. |
| 53281018 | Failover policy is not compatible with one or more services in service policy |
| 53281065 | The service_policy does not exist in the SVM. |
| 53281072 | The failover policy is not valid for interfaces in the SVM. |
| 53281086 | LIF would exceed the maximum number of supported intercluster LIFs in IPspace. |
| 53281087 | SVM must have an NVMe service to perform this operation. |
| 53281089 | LIF on SVM cannot be updated to use service policy because that service policy includes SAN services and the target LIF is not home. |

| Error Code | Description |
| --- | --- |
| 53281096 | LIF could not be modified because the current port does not support the specified rdma-protocols. |
| 53281106 | Failed checking the cluster capabilities. |
| 53281109 | The interface could not be updated to use the service policy because the interface is currently associated with BGP peer group. |
| 53281111 | Cannot update the service policy because the interface is associated with a BGP peer-group. |
| 53281113 | Configuration details for LIF were modified but an error was encountered while updating the LIF's operational status. |
| 53281114 | The specified parameter is only supported on data SVMs. |
| 53281468 | Failed to update external route tables for IP. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
|---|---|---|
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

home_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

location

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port.

| Name | Type | Description |
| --- | --- | --- |
| auto_revert | boolean | |

| Name | Type | Description |
| --- | --- | --- |
| failover | string | Policy that defines where an interface is permitted to move on failover. The policy named “default” implements the recommended best practice for NAS LIFs on the current platform and cluster, and was known as “system_defined” in the CLI. |
| home_node | home_node | |
| home_port | home_port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |
| is_home | boolean | |
| node | node | |
| port | port | Port UUID along with readable names. Either the UUID or both names may be supplied on input. |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

The most recent sample of I/O metrics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |

service_policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

throughput_raw

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

The real time I/O statistics for the interface.

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the subnet. If only the name is provided, the IPspace scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_interface

| Name | Type | Description |
|---|---|---|
| ddns_enabled | boolean | Indicates whether or not dynamic DNS updates are enabled. Defaults to true if the interface supports "data_nfs" or "data_cifs" services, otherwise false. |
| dns_zone | string | Fully qualified DNS zone name |
| enabled | boolean | The administrative state of the interface. |

| Name | Type | Description |
| --- | --- | --- |
| fail_if_subnet_conflicts | boolean | This command fails if the specified IP address falls within the address range of a named subnet. Set this value to false to use the specified IP address and to assign the subnet owning that address to the interface. |
| ip | ip_info | IP information |
| location | location | Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node. For PATCH, set is_home to true to revert a LIF back to its home port. |
| name | string | Interface name |
| rdma_protocols | array[string] | Supported RDMA offload protocols |
| service_policy | service_policy | |
| services | array[string] | The services associated with the interface. |
| state | string | The operational state of the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve interface historical performance metrics**

GET `/network/ip/interfaces/{uuid}/metrics`

**Introduced In:** 9.8

Retrieves historical performance metrics for an interface.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| uuid | string | path | True | Unique identifier of the interface. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

throughput

The rate of throughput bytes per second observed at the interface.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Throughput performance for the interfaces.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the interface. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The UUID that uniquely identifies the interface. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network IP routes

### Manage network IP routes

#### Overview

This endpoint supports the following operations: GET (collection and instance), POST, and DELETE.

---

#### Retrieving network routes

You can use the IP routes GET API to retrieve and display relevant information pertaining to the routes configured in the cluster. The API retrieves the list of all routes configured in the cluster, or a specific route. The fields that are returned in the response will differ with the configuration.

#### Examples

### Retrieving all routes in the cluster

The following output shows the list of all routes configured in a cluster.

---

```
# The API:
/api/network/ip/routes

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/routes?fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "5fdffb0b-62f8-11e8-853d-005056b4c971",
    "ipspace": {
      "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
      "name": "Default",
```

```
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-
005056b4c971"
        }
      }
    },
    "svm": {
      "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
        }
      }
    },
    "scope": "svm",
    "destination": {
      "address": "10.4.3.14",
      "netmask": "18",
      "family": "ipv4"
    },
    "gateway": "10.4.3.1",
    "_links": {
      "self": {
        "href": "/api/network/ip/routes/5fdffb0b-62f8-11e8-853d-
005056b4c971"
      }
    }
  },
  {
    "uuid": "84c128d2-62f9-11e8-853d-005056b4c971",
    "ipspace": {
      "uuid": "cc71aadc-62f7-11e8-853d-005056b4c971",
      "name": "ips1",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/cc71aadc-62f7-11e8-853d-
005056b4c971"
        }
      }
    },
    "scope": "cluster",
    "destination": {
      "address": "::",
      "netmask": "0",
```

```
      "family": "ipv6"
    },
    "gateway": "fd20:8b1e:b255:814e::1",
    "_links": {
      "self": {
        "href": "/api/network/ip/routes/84c128d2-62f9-11e8-853d-
005056b4c971"
      }
    }
  },
  {
    "uuid": "8cc72bcd-616c-11e8-a4df-005056b4c971",
    "ipspace": {
      "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-
005056b4c971"
        }
      }
    },
    "scope": "cluster",
    "destination": {
      "address": "0.0.0.0",
      "netmask": "0",
      "family": "ipv4"
    },
    "gateway": "10.224.64.1",
    "_links": {
      "self": {
        "href": "/api/network/ip/routes/8cc72bcd-616c-11e8-a4df-
005056b4c971"
      }
    }
  },
  {
    "uuid": "d63b6eee-62f9-11e8-853d-005056b4c971",
    "ipspace": {
      "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-
005056b4c971"
        }
```

```
      }
    },
    "svm": {
      "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
        }
      }
    },
    "scope": "svm",
    "destination": {
      "address": "fd20:8b1e:b255:814e::",
      "netmask": "64",
      "family": "ipv6"
    },
    "gateway": "fd20:8b1e:b255:814e::1",
    "_links": {
      "self": {
        "href": "/api/network/ip/routes/d63b6eee-62f9-11e8-853d-
005056b4c971"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/network/ip/routes?fields=*"
  }
}
}
```


### Retrieving a specific Cluster-scoped route

The following output shows the returned response when a specific Cluster-scoped route is requested. The system returns an error if there is no route with the requested UUID. SVM information is not returned for Cluster-scoped routes.

```
# The API:
/api/network/ip/routes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/routes/84c128d2-62f9-11e8-
853d-005056b4c971?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "84c128d2-62f9-11e8-853d-005056b4c971",
"ipspace": {
  "uuid": "cc71aadc-62f7-11e8-853d-005056b4c971",
  "name": "ips1",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/cc71aadc-62f7-11e8-853d-005056b4c971"
    }
  }
},
"scope": "cluster",
"destination": {
  "address": "::",
  "netmask": "0",
  "family": "ipv6"
},
"gateway": "fd20:8b1e:b255:814e::1",
"_links": {
  "self": {
    "href": "/api/network/ip/routes/84c128d2-62f9-11e8-853d-005056b4c971"
  }
}
}
```

#### Retrieving a specific SVM-scoped route

The following output shows the returned response when a specific SVM-scoped route is requested. The system returns an error if there is no route with the requested UUID. The SVM object is only included for SVM-scoped routes.

```
# The API:
/api/network/ip/routes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/routes/d63b6eee-62f9-11e8-
853d-005056b4c971?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "d63b6eee-62f9-11e8-853d-005056b4c971",
"ipspace": {
  "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-005056b4c971"
    }
  }
},
"svm": {
  "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
    }
  }
},
"scope": "svm",
"destination": {
  "address": "fd20:8b1e:b255:814e::",
  "netmask": "64",
  "family": "ipv6"
},
"gateway": "fd20:8b1e:b255:814e::1",
"_links": {
  "self": {
    "href": "/api/network/ip/routes/d63b6eee-62f9-11e8-853d-005056b4c971"
  }
}
}
```

#### Creating network routes

You can use the POST API to create an SVM-scoped route by specifying the associated SVM, or a Cluster-scoped route by specifying the associated IPspace.

#### Examples

### Creating a Cluster-scoped route

IPspace is required to create a Cluster-scoped route. If the IPspace is not specified, the route will be created in the Default IPspace. The default destination will be set to "0.0.0.0/0" for IPv4 gateway addresses or "::/0" for IPv6 gateway addresses.

```
# The API:
/api/network/ip/routes

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/routes?return_records=true"
-H "accept: application/json" -d '{ "ipspace": { "name":"ips1" },
"gateway": "10.10.10.1"}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "ae583c9e-9ac7-11e8-8bc9-005056bbd531",
    "ipspace": {
      "name": "ips1"
    },
    "gateway": "10.10.10.1"
  }
]
}
```

### Creating an SVM-scoped route

To create an SVM-scoped route, the associated SVM can be identified by either its UUID or name.

```
# The API:
/api/network/ip/routes

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/routes?return_records=true"
-H "accept: application/json" -d '{ "svm": { "name":"vs0" }, "gateway":
"10.10.10.1"}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "38805a91-9ac9-11e8-8bc9-005056bbd531",
    "svm": {
      "name": "vs0"
    },
    "gateway": "10.10.10.1"
  }
]
}
```

---

#### Deleting network routes

You can use the DELETE API to delete a specific route identified by its UUID.

#### Example

### Deleting a specific route

---

```
# The API:
/api/network/ip/routes/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/routes/38805a91-9ac9-
11e8-8bc9-005056bbd531"
```

---

### Retrieve IP routes

GET `/network/ip/routes`

**Introduced In:** 9.6

Retrieves the collection of IP routes.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `interfaces.*`

#### Related ONTAP commands

- `network route show`
- `network route show-lifs`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interfaces.ip.address | string | query | False | Filter by interfaces.ip.address<br>• Introduced in: 9.9 |
| interfaces.name | string | query | False | Filter by interfaces.name<br>• Introduced in: 9.9 |
| interfaces.uuid | string | query | False | Filter by interfaces.uuid<br>• Introduced in: 9.9 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| scope | string | query | False | Filter by scope |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | query | False | Filter by uuid |
| destination.family | string | query | False | Filter by destination.family |
| destination.address | string | query | False | Filter by destination.address |
| destination.netmask | string | query | False | Filter by destination.netmask |
| gateway | string | query | False | Filter by gateway |
| metric | integer | query | False | Filter by metric<br>• Introduced in: 9.11 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| error | error | |
| num_records | integer | Number of records |
| records | array[network_route] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "destination": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "gateway": "10.1.1.1",
      "interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
```

```
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "ipspace": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "Default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "scope": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

network_route

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |
| interfaces | array[interfaces] | IP interfaces on the same subnet as the gateway. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input. |
| metric | integer | Indicates a preference order between several routes to the same destination. With typical usage, the default metrics provided are adequate, there is no need to specify a metric in the route creation. |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied. |
| uuid | string | The UUID that uniquely identifies the route. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a cluster-scoped or SVM-scoped static route

POST `/network/ip/routes`

**Introduced In:** 9.6

Creates a Cluster-scoped or SVM-scoped static route.

#### Required properties

- `gateway` - IP address to route packets to.
- SVM-scoped routes
  - `svm.name` or `svm.uuid` - SVM that route is applied to.
- cluster-scoped routes
  - There are no additional required fields for Cluster-scoped routes.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `destination` - *0.0.0.0/0* for IPv4 or *::/0* for IPv6.
- `ipspace.name`
  - *Default* for Cluster-scoped routes.
  - Name of the SVM's IPspace for SVM-scoped routes.
- `metric` - 20.

#### Related ONTAP commands

- `network route create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |
| interfaces | array[interfaces] | IP interfaces on the same subnet as the gateway. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input. |
| metric | integer | Indicates a preference order between several routes to the same destination. With typical usage, the default metrics provided are adequate, there is no need to specify a metric in the route creation. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied. |
| uuid | string | The UUID that uniquely identifies the route. |

### Example request

```
{
  "destination": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[network_route] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "destination": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "gateway": "10.1.1.1",
      "interfaces": [
        {
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1966265 | The destination and gateway must belong to the same IP address family. |
| 1966345 | Duplicate route exists. |
| 1966505 | The gateway address cannot be the same as a LIF address. A LIF is already configured with that IP address. |
| 1967080 | The destination.address is missing. |
| 1967081 | The specified SVM must exist in the specified IPspace. |
| 1967082 | The specified ipspace.uuid and ipspace.name refer to different IPspaces. |
| 1967146 | The specified svm.name is not valid. |
| 2621462 | The specified SVM does not exist. |
| 2621574 | This operation is not permitted on an SVM that is configured as the destination of a MetroCluster SVM relationship. |
| 2621706 | The specified UUID and name refer to different SVMs. |
| 53282375 | The specified destination is invalid. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip

IP information

interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

network_route

| Name | Type | Description |
| --- | --- | --- |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |
| interfaces | array[interfaces] | IP interfaces on the same subnet as the gateway. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input. |
| metric | integer | Indicates a preference order between several routes to the same destination. With typical usage, the default metrics provided are adequate, there is no need to specify a metric in the route creation. |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the route. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an IP route

DELETE `/network/ip/routes/{uuid}`

**Introduced In:** 9.6

Deletes a specific IP route.

#### Related ONTAP commands

- `network route delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Route UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve details for an IP route

GET `/network/ip/routes/{uuid}`

**Introduced In:** 9.6

Retrieves the details of a specific IP route.

#### Related ONTAP commands

- `network route show`
- `network route show-lifs`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Route UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |
| interfaces | array[interfaces] | IP interfaces on the same subnet as the gateway. |
| ipspace | ipspace | Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input. |
| metric | integer | Indicates a preference order between several routes to the same destination. With typical usage, the default metrics provided are adequate, there is no need to specify a metric in the route creation. |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| svm | svm | Applies only to SVM-scoped objects. Either the UUID or name must be supplied. |
| uuid | string | The UUID that uniquely identifies the route. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "destination": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interfaces

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network IP service policies

### Manage network IP service policies

#### Overview

Service policies are named groupings that define what services are supported by an IP interface.
The following operations are supported:

- Creation: POST network/ip/service-policies
- Collection Get: GET network/ip/service-policies
- Instance Get: GET network/ip/service-policies/{uuid}
- Instance Patch: PATCH network/ip/service-policies/{uuid}
- Instance Delete: DELETE network/ip/service-polices/{uuid}

#### Examples

### Retrieving all service policies in the cluster

The following output shows the collection of all service policies configured in a 2-node cluster. By default (without 'field=*' parameter), only the UUID and name fields are shown for each entry.

```
# The API:
/api/network/ethernet/ip/service-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
```

```
    "uuid": "e4e2f193-c1a3-11e8-bb9d-005056bb88c8",
    "name": "net-intercluster",
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/e4e2f193-c1a3-11e8-bb9d-
005056bb88c8"
      }
    }
  },
  {
    "uuid": "e4e3f6da-c1a3-11e8-bb9d-005056bb88c8",
    "name": "net-route-announce",
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/e4e3f6da-c1a3-11e8-bb9d-
005056bb88c8"
      }
    }
  },
  {
    "uuid": "e5111111-c1a3-11e8-bb9d-005056bb88c8",
    "name": "vserver-route-announce",
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/e5111111-c1a3-11e8-bb9d-
005056bb88c8"
      }
    }
  },
  {
    "uuid": "e6111111-c1a3-11e8-bb9d-005056bb88c8",
    "name": "data-route-announce",
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/e6111111-c1a3-11e8-bb9d-
005056bb88c8"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/network/ip/service-
policies/?return_records=true&return_timeout=15"
  }
```

```
  }
}
```

### Retrieving a specific service policy (scope=svm)

The following output displays the response when a specific "svm" scoped service policy is requested. Among other parameters, the response contains the svm parameters associated with the service policy. The system returns an error when there is no service policy with the requested UUID.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "http://<mgmt-ip>/api/network/ip/service-policies/dad323ff-
4ce0-11e9-9372-005056bb91a8?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "dad323ff-4ce0-11e9-9372-005056bb91a8",
"name": "default-data-files",
"scope": "svm",
"svm": {
  "uuid": "d9060680-4ce0-11e9-9372-005056bb91a8",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/d9060680-4ce0-11e9-9372-005056bb91a8"
    }
  }
},
"ipspace": {
  "uuid": "45ec2dee-4ce0-11e9-9372-005056bb91a8",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/45ec2dee-4ce0-11e9-9372-005056bb91a8"
    }
  }
},
"is_built_in": true,
"services": [
  "data_core",
  "data_nfs",
  "data_cifs",
  "data_flexcache"
],
"_links": {
  "self": {
    "href": "/api/network/ip/service-policies/dad323ff-4ce0-11e9-9372-
005056bb91a8"
  }
}
}
```

#### Retrieving a specific service policy (scope=svm) when requesting commonly used fields

The following output displays the response when commonly used fields are requested for a specific "svm" scoped service policy. Among other parameters, the response contains the svm parameters associated with the service policy. The system returns an error when there is no service policy with the requested UUID.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/e0889ce6-
1e6a-11e9-89d6-005056bbdc04?fields=name,scope,svm.name,ipspace.name" -H
"accept: application/hal+json"

# The response:
{
"uuid": "e0889ce6-1e6a-11e9-89d6-005056bbdc04",
"name": "test_policy",
"scope": "svm",
"svm": {
  "name": "vs0"
},
"ipspace": {
  "name": "Default"
},
"_links": {
  "self": {
    "href": "/api/network/ip/service-policies/e0889ce6-1e6a-11e9-89d6-
005056bbdc04"
  }
}
}
```

#### Retrieving a specific service policy (scope=cluster)

The following output displays the response when a specific cluster-scoped service policy is requested. The SVM object is not included for cluster-scoped service policies. A service policy with a scope of "cluster" is associated with an IPspace. The system returns an error when there is no service policy with the requested UUID.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/4c6b72b9-
0f6c-11e9-875d-005056bb21b8?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "4c6b72b9-0f6c-11e9-875d-005056bb21b8",
"name": "net-intercluster",
"scope": "cluster",
"ipspace": {
  "uuid": "4051f13e-0f6c-11e9-875d-005056bb21b8",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/4051f13e-0f6c-11e9-875d-005056bb21b8"
    }
  }
},
"is_built_in": false,
"services": [
  "intercluster_core"
],
"_links": {
  "self": {
    "href": "/api/network/ip/service-policies/4c6b72b9-0f6c-11e9-875d-
005056bb21b8"
  }
}
}
```

#### Retrieving a specific service policy (scope=cluster) when requesting commonly used fields

The following output displays the response when commonly used fields are requested for a specific "cluster" scoped service policy. The SVM object is not included for cluster-scoped service policies. A service policy with a scope of "cluster" is associated with an IPspace. The system returns an error when there is no service policy with the requested UUID.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/4c6b72b9-
0f6c-11e9-875d-005056bb21b8?fields=name,scope,ipspace.name" -H "accept:
application/hal+json"

# The response:
{
"uuid": "4c6b72b9-0f6c-11e9-875d-005056bb21b8",
"name": "net-intercluster",
"scope": "cluster",
"ipspace": {
  "name": "Default"
},
"services": [
  "intercluster_core"
],
"_links": {
  "self": {
    "href": "/api/network/ip/service-policies/4c6b72b9-0f6c-11e9-875d-
005056bb21b8"
  }
}
}
```

---

#### Creating service policies

You can use this API to create an SVM-scoped service policy by specifying the associated SVM, or a cluster-scoped service policy by specifying the associated IPspace. If the scope is not specified, it is inferred from the presence of the IPspace or SVM.
Cluster scoped service policies will operate on the IPspace "Default" unless IPspace is explicitly specified.

#### Examples

### Creating a cluster-scoped service policy

The following output displays the response when creating a service policy with a scope of "cluster" and an IPspace of "Default".

---

```
# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "scope": "cluster", "ipspace": { "name":"Default" },
"services": [ "intercluster_core" ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "74139267-f1aa-11e9-b5d7-005056a73e2e",
    "name": "new-policy",
    "scope": "cluster",
    "ipspace": {
      "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-
005056a7080e"
        }
      }
    },
    "is_built_in": false,
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-
005056a73e2e"
      }
    },
    "services": [
      "intercluster_core"
    ]
  }
]
}
```

#### Creating a cluster-scoped service policy without specifying IPspace

The following output displays the response when creating a service policy with a scope of "cluster" without specifying an IPspace".

```
# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "scope": "cluster", "services": [ "intercluster_core" ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "74139267-f1aa-11e9-b5d7-005056a73e2e",
    "name": "new-policy",
    "scope": "cluster",
    "ipspace": {
      "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-
005056a7080e"
        }
      }
    },
    "is_built_in": false,
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-
005056a73e2e"
      }
    },
    "services": [
      "intercluster_core"
    ]
  }
]
}
```

### Creating a cluster-scoped service policy without specifying scope

The following output displays the response when creating a service policy in the "Default" IPspace without specifying the scope".

```
# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy2", "ipspace.name": "Default", "services": [
"intercluster_core" ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "59439267-f1aa-11e9-b5d7-005056a73e2e",
    "name": "new-policy2",
    "scope": "cluster",
    "ipspace": {
      "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-
005056a7080e"
        }
      }
    },
    "is_built_in": false,
    "services": [
      "intercluster_core"
    ],
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-
005056a73e2e"
      }
    }
  }
]
}
```

### Creating an SVM-scoped service policy

The following output displays the response when creating a service policy with a scope of "svm" in the SVM "vs0".

```
# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "scope": "svm", "svm": { "name":"vs0" }, "services": [
"data-nfs", "data-cifs" ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "f3901097-f2c4-11e9-b5d7-005056a73e2e",
    "name": "new-policy",
    "scope": "svm",
    "svm": {
      "uuid": "07df9cee-e912-11e9-a13a-005056a73e2e",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/07df9cee-e912-11e9-a13a-005056a73e2e"
        }
      }
    },
    "ipspace": {
      "uuid": "1d3199d2-e906-11e9-a13a-005056a73e2e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/1d3199d2-e906-11e9-a13a-
005056a73e2e"
        }
      }
    },
    "is_built_in": false,
    "services": [
      "data_nfs",
      "data_cifs"
```

```
    ],
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/f3901097-f2c4-11e9-b5d7-
005056a73e2e"
      }
    }
  }
]
}
```

#### Creating an SVM-scoped service policy without specifying scope

The following output displays the response when creating a service policy with a SVM of "vs0" without specifying the scope.

```
# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "svm": { "name":"vs0" }, "services": [ "data-nfs", "data-
cifs" ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "f3901097-f2c4-11e9-b5d7-005056a73e2e",
    "name": "new-policy",
    "scope": "svm",
    "svm": {
      "uuid": "07df9cee-e912-11e9-a13a-005056a73e2e",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/07df9cee-e912-11e9-a13a-005056a73e2e"
        }
      }
    },
    "ipspace": {
```

```
      "uuid": "1d3199d2-e906-11e9-a13a-005056a73e2e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/1d3199d2-e906-11e9-a13a-
005056a73e2e"
        }
      }
    },
    "is_built_in": false,
    "services": [
      "data_nfs",
      "data_cifs"
    ],
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/f3901097-f2c4-11e9-b5d7-
005056a73e2e"
      }
    }
  }
]
}
```


#### Updating the name of a service policy

The following example displays the command used to update the name of a service policy scoped to a specific "svm". The system returns an error when there is no
service policy associated with the UUID or the service policy cannot be renamed.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/service-policies/734eaf57-
d2fe-11e9-9284-005056acaad4" -d "{ \"name\": \"new-name\" }" -H "accept:
application/hal+json"
```

#### Updating the services for a service policy

The following example displays the command used to update the services a service policy contains.

The specified services replace the existing services. To retain existing services, they must be included in the PATCH request.
The system returns an error when there is no
service policy associated with the UUID or the services cannot be applied.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/service-policies/734eaf57-
d2fe-11e9-9284-005056acaad4" -d "{ \"services\": [ \"data-nfs\",\"data-
cifs\" ] }" -H "accept: application/hal+json"
```

### Deleting a service policy

The following output displays the response for deleting a service policy.

```
# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X  DELETE "https://<mgmt-ip>/api/network/ip/service-
policies/757ed726-bdc1-11e9-8a92-005056a7bf25" -H "accept:
application/hal+json"
```

### Retrieve service policies

GET `/network/ip/service-policies`

**Introduced In:** 9.6

Retrieves a collection of service policies.

#### Related ONTAP commands

- `network interface service-policy show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name |
| services | string | query | False | Filter by services |
| scope | string | query | False | Filter by scope |
| is_built_in | boolean | query | False | Filter by is_built_in<br>• Introduced in: 9.11 |
| uuid | string | query | False | Filter by uuid |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[ip_service_policy] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "default-intercluster",
      "scope": "string",
      "services": [
        "data_nfs"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_service_policy

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | |
| is_built_in | boolean | |
| name | string | |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| services | array[string] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a service policy for network interfaces**

POST `/network/ip/service-policies`

**Introduced In:** 9.8

Creates a service policy for network interfaces.

#### Related ONTAP commands

- `network interface service-policy create`

#### Required properties

- `name` - Name of the service policy to create.
- `ipspace.name` or `ipspace.uuid`
  - Required for cluster-scoped service policies.
  - Optional for SVM-scoped service policies.
- `svm.name` or `svm.uuid`
  - Required for SVM-scoped service policies.
  - Not valid for cluster-scoped service policies.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `scope`
  - svm if the svm parameter is specified
  - cluster if the svm parameter is not specified

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | |
| is_built_in | boolean | |
| name | string | |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| services | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

#### Example request

```
{
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "default-intercluster",
  "scope": "string",
  "services": [
    "data_nfs"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1966373 | Port must reside in the same IPspace as the interface's SVM. |
| 1967146 | Svm.name does not exist. |
| 1967147 | Svm.uuid does not exist. |
| 2621740 | An unexpected error when trying to determine whether the target Vserver was locked or not on this cluster. |
| 53281895 | Custom service policies cannot be created on the specified svm |
| 53281929 | Service policies cannot combine block and file services. |
| 53281931 | Service policy names cannot start with "default-". |
| 53281932 | Service cannot be added because the service does not exist for the specified SVM or IPspace. |
| 53281933 | A Cluster-scoped service cannot be added to a SVM-scoped service policy. |
| 53281934 | An SVM-scoped service cannot be added to a Cluster-scoped service policy. |
| 53281935 | Scope is set to "svm" and svm.uuid or svm.name have not been specified. |
| 53281936 | The SVM is not in the specified IPspace. |
| 53281937 | Svm.uuid and svm.name are not valid parameters when scope is cluster. |
| 53281938 | Svm.uuid or svm.name specify a vserver that does not exist. |
| 53281939 | One or more of the svm.uuid, svm.name, ipspace.uuid, and ipspace.name have invalid values. |
| 53281940 | SVM or IPspace has not been specified. |
| 53281941 | SVM does not exist. |
| 53281942 | Service policy cannot be created because the specified SVM does not exist. |
| 53281944 | Ipspace.name does not exist. |
| 53281945 | Ipspace.uuid is not an IPspace. |
| 53281946 | Service policy already exists. |
| 53281958 | Service policies cannot contain multiple block-oriented services. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_service_policy

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | |
| is_built_in | boolean | |
| name | string | |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| services | array[string] | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a service policy for network interfaces

DELETE `/network/ip/service-policies/{uuid}`

**Introduced In:** 9.8

Deletes a service policy for network interfaces.

#### Related ONTAP commands

- `network interface service-policy delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the service policy |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621740 | An unexpected error when trying to determine whether the target Vserver was locked or not on this cluster. |
| 53281927 | Service policies owned by the system cannot be deleted. |
| 53281928 | Service policies assigned to LIFs cannot be deleted. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a service policy

GET `/network/ip/service-policies/{uuid}`

**Introduced In:** 9.6

Retrieves a specific service policy.

#### Related ONTAP commands

- `network interface service-policy show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Service policy UUID |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ipspace | ipspace | |
| is_built_in | boolean | |
| name | string | |
| scope | string | Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster". |
| services | array[string] | |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "default-intercluster",
  "scope": "string",
  "services": [
    "data_nfs"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a service policy for network interfaces

PATCH `/network/ip/service-policies/{uuid}`

**Introduced In:** 9.8

Updates a service policy for network interfaces.

#### Related ONTAP commands

- `network interface service-policy add-service`
- `network interface service-policy modify-service`
- `network interface service-policy remove-service`
- `network interface service-policy rename`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | UUID of the service policy |

#### Request Body

| Name | Type | Description |
|---|---|---|
| is_built_in | boolean | |
| name | string | |
| services | array[string] | |
| uuid | string | |

#### Example request

```
{
  "name": "default-intercluster",
  "services": [
    "data_nfs"
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1376669 | Port must reside in the same IPspace as the interface's SVM. |
| 2621740 | An unexpected error when trying to determine whether the target Vserver was locked or not on this cluster. |
| 53281911 | Modifying the service policy is disallowed because policies on this SVM are managed by the system |
| 53281929 | Service policies cannot combine block and file services. |
| 53281930 | Service policies maintained by the system cannot be renamed. |
| 53281931 | Service policy names cannot start with "default-". |
| 53281932 | Service cannot be added because the service does not exist for the specified SVM or IPspace. |
| 53281933 | A Cluster-scoped service cannot be added to a SVM-scoped service policy. |
| 53281934 | An SVM-scoped service cannot be added to a Cluster-scoped service policy. |

| Error Code | Description |
| --- | --- |
| 53281952 | The service policy on an SVM cannot be updated to include a block service. Use built-in service policies for SAN services. |
| 53281953 | The service policy on an SVM cannot be updated to include a new service. |
| 53281960 | Service cannot be removed from the service policy because it is used by one or more interfaces. |
| 53281961 | Service cannot be removed from the service policy because it is used by one or more interfaces. |
| 53281963 | Service cannot be removed from the service policy because it is used by one or more interfaces. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ip_service_policy

| Name | Type | Description |
|---|---|---|
| is_built_in | boolean | |
| name | string | |
| services | array[string] | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network IP subnets

### Manage network IP subnets

#### Overview

The following operations are supported:

- Creation: POST network/ip/subnets
- Collection Get: GET network/ip/subnets
- Instance Get: GET network/ip/subnets/{uuid}
- Instance Patch: PATCH network/ip/subnets/{uuid}
- Instance Delete: DELETE network/ip/subnets/{uuid}

#### Retrieving IP subnet information

The IP subnets GET API retrieves and displays relevant information pertaining to the subnets configured in the cluster. The response can contain a list of multiple subnets or a specific subnet.

#### Examples

### Retrieving all subnets in the cluster

The following example shows the list of all subnets configured in a cluster.

```
# The API:
/api/network/ip/subnets

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "451d8d99-582c-11ec-8572-005056acd597",
    "name": "Subnet-002",
    "_links": {
      "self": {
        "href": "/api/network/ip/subnets/451d8d99-582c-11ec-8572-
005056acd597"
      }
    }
  },
  {
    "uuid": "615b722f-5795-11ec-8572-005056acd597",
    "name": "Subnet-001",
    "_links": {
      "self": {
        "href": "/api/network/ip/subnets/615b722f-5795-11ec-8572-
005056acd597"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/ip/subnets"
  }
}
}
```

#### Retrieving a specific subnet

The following example shows the response when a specific subnet is requested. This is equivalent to fields=*, which returns most of the fields. The system returns an error when there is no subnet with the requested UUID.

```
# The API:
/api/network/ip/subnets/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets/451d8d99-582c-11ec-
8572-005056acd597" -H "accept: application/hal+json"

# The response:
{
"uuid": "451d8d99-582c-11ec-8572-005056acd597",
"name": "Subnet-002",
"ipspace": {
  "uuid": "6f62c691-5780-11ec-8572-005056acd597",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/6f62c691-5780-11ec-8572-005056acd597"
    }
  }
},
"broadcast_domain": {
  "uuid": "9a1dce3b-5780-11ec-8572-005056acd597",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/9a1dce3b-5780-11ec-
8572-005056acd597"
    }
  }
},
"subnet": {
  "address": "10.2.1.0",
  "netmask": "24",
  "family": "ipv4"
},
"gateway": "10.2.1.1",
"_links": {
  "self": {
    "href": "/api/network/ip/subnets/451d8d99-582c-11ec-8572-005056acd597"
  }
}
}
```

#### Retrieving all the fields for a specific subnet

The following example shows the response when all the fields for a specific subnet are requested, returning everything that fields=* returns plus the IP ranges and count fields. The system returns an error when there is no subnet with the requested UUID.

```
# The API:
/api/network/ip/subnets/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets/451d8d99-582c-11ec-
8572-005056acd597?fields=**" -H "accept: application/hal+json"

# The response:
{
"uuid": "451d8d99-582c-11ec-8572-005056acd597",
"name": "Subnet-002",
"ipspace": {
  "uuid": "6f62c691-5780-11ec-8572-005056acd597",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/6f62c691-5780-11ec-8572-005056acd597"
    }
  }
},
"broadcast_domain": {
  "uuid": "9a1dce3b-5780-11ec-8572-005056acd597",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/9a1dce3b-5780-11ec-
8572-005056acd597"
    }
  }
},
"subnet": {
  "address": "10.2.1.0",
  "netmask": "24",
  "family": "ipv4"
},
"gateway": "10.2.1.1",
"ip_ranges": [
  {
    "start": "10.2.1.10",
```

```
    "end": "10.2.1.22",
    "family": "ipv4"
  },
  {
    "start": "10.2.1.101",
    "end": "10.2.1.200",
    "family": "ipv4"
  }
],
"available_ip_ranges": [
  {
    "start": "10.2.1.10",
    "end": "10.2.1.22",
    "family": "ipv4"
  },
  {
    "start": "10.2.1.101",
    "end": "10.2.1.200",
    "family": "ipv4"
  }
],
"total_count": 113,
"used_count": 0,
"available_count": 113,
"_links": {
  "self": {
    "href": "/api/network/ip/subnets/451d8d99-582c-11ec-8572-
005056acd597?fields=**"
  }
}
}
```


#### Creating IP subnets

You can use the IP subnets POST API to create IP subnets as shown in the following examples.

#### Examples

**Creating an IP subnet using the minimum number of parameters.**

The following example shows the record returned after the creation of an IP subnet.

```
# The API:
/api/network/ip/subnets

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/subnets" -H "accept:
application/hal+json" -d '{ "name": "Subnet-003", "broadcast_domain": {
"uuid": "6577524b-5863-11ec-8981-005056a7077f" }, "subnet": { "address":
"10.3.0.0", "netmask": "16" } }'

# The response:
{
}

# The result:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets?name=Subnet-
003&fields=**" -H "accept: application/hal+json"
{
"records": [
  {
    "uuid": "79ff5a5e-59b7-11ec-8981-005056a7077f",
    "name": "Subnet-003",
    "ipspace": {
      "uuid": "36569d0f-5863-11ec-8981-005056a7077f",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/36569d0f-5863-11ec-8981-
005056a7077f"
        }
      }
    },
    "broadcast_domain": {
      "uuid": "6577524b-5863-11ec-8981-005056a7077f",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/6577524b-5863-
11ec-8981-005056a7077f"
        }
      }
    },
    "subnet": {
      "address": "10.3.0.0",
      "netmask": "16",
      "family": "ipv4"
    },
```

```
    "total_count": 0,
    "used_count": 0,
    "available_count": 0,
    "_links": {
      "self": {
        "href": "/api/network/ip/subnets/79ff5a5e-59b7-11ec-8981-
005056a7077f?fields=**"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/subnets?name=Subnet-003&fields=**"
  }
}
}
```

**Creating an IP subnet using all parameters.**

The following example shows the record returned after the creation of an IP subnet setting all parameters.

```
# The API:
/api/network/ip/subnets

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/subnets" -H "accept:
application/hal+json" -d '{ "name": "Subnet-004", "ipspace": { "name":
"Default", "uuid": "36569d0f-5863-11ec-8981-005056a7077f" },
"broadcast_domain": { "name": "Default", "uuid": "6577524b-5863-11ec-8981-
005056a7077f" }, "subnet": { "address": "10.4.1.0", "netmask": "24" },
"gateway": "10.4.1.1", "ip_ranges": [ { "start": "10.4.1.30", "end":
"10.4.1.39" }, { "start": "10.4.1.150", "end": "10.4.1.229" } ],
"fail_if_lifs_conflict": "false" }'

# The response:
{
}

# The result:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets?name=Subnet-
```

```
004&fields=**" -H "accept: application/hal+json"
{
"records": [
  {
    "uuid": "0e0a19e7-59ba-11ec-8981-005056a7077f",
    "name": "Subnet-004",
    "ipspace": {
      "uuid": "36569d0f-5863-11ec-8981-005056a7077f",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/36569d0f-5863-11ec-8981-
005056a7077f"
        }
      }
    },
    "broadcast_domain": {
      "uuid": "6577524b-5863-11ec-8981-005056a7077f",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/6577524b-5863-
11ec-8981-005056a7077f"
        }
      }
    },
    "subnet": {
      "address": "10.4.1.0",
      "netmask": "24",
      "family": "ipv4"
    },
    "gateway": "10.4.1.1",
    "ip_ranges": [
      {
        "start": "10.4.1.30",
        "end": "10.4.1.39",
        "family": "ipv4"
      },
      {
        "start": "10.4.1.150",
        "end": "10.4.1.229",
        "family": "ipv4"
      }
    ],
    "available_ip_ranges": [
      {
```

```
        "start": "10.4.1.30",
        "end": "10.4.1.39",
        "family": "ipv4"
      },
      {
        "start": "10.4.1.150",
        "end": "10.4.1.229",
        "family": "ipv4"
      }
    ],
    "total_count": 90,
    "used_count": 0,
    "available_count": 90,
    "_links": {
      "self": {
        "href": "/api/network/ip/subnets/0e0a19e7-59ba-11ec-8981-
005056a7077f?fields=**"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/subnets?name=Subnet-004&fields=**"
  }
}
}
```


#### Updating IP subnets

You can use the IP subnets PATCH API to update the attributes of an IP subnet.

#### Examples

### Updating the name of an IP subnet

The following example shows how the PATCH request changes the name.

```
# The API:
/api/network/ip/subnets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-
11ec-8981-005056a7077f" -H "accept: application/hal+json" -d '{ "name":
"Subnet-004-NewName" }'

# The response:
{
}

# The result:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-11ec-
8981-005056a7077f?fields=name" -H "accept: application/hal+json"
{
"uuid": "0e0a19e7-59ba-11ec-8981-005056a7077f",
"name": "Subnet-004-NewName",
"_links": {
  "self": {
    "href": "/api/network/ip/subnets/0e0a19e7-59ba-11ec-8981-005056a7077f"
  }
}
}
```

#### Updating the ip_ranges of an IP subnet

The following example shows how the PATCH request updates the ip_ranges.

```
# The API:
/api/network/ip/subnets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-
11ec-8981-005056a7077f" -H "accept: application/hal+json" -d '{
"ip_ranges": [ { "start": "10.4.1.20", "end": "10.4.1.239" } ] }'

# The response:
{
}

# The result:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-11ec-
8981-005056a7077f?fields=ip_ranges" -H "accept: application/hal+json"
{
"uuid": "0e0a19e7-59ba-11ec-8981-005056a7077f",
"name": "Subnet-004-NewName",
"ip_ranges": [
  {
    "start": "10.4.1.20",
    "end": "10.4.1.239",
    "family": "ipv4"
  }
],
"_links": {
  "self": {
    "href": "/api/network/ip/subnets/0e0a19e7-59ba-11ec-8981-005056a7077f"
  }
}
}
```

#### Deleting IP subnets

You can use the IP subnets DELETE API to delete an IP subnet.

#### Example

### Deleting an IP subnet

The following DELETE request deletes a specific network IP subnet.

```
# The API:
/api/network/ip/subnets/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-
11ec-8981-005056a7077f"

# The response:
{
}

# The result:
curl -X GET "https://<mgmt-ip>/api/network/ip/subnets/0e0a19e7-59ba-11ec-
8981-005056a7077f" -H "accept: application/hal+json"
{
"error": {
  "message": "entry doesn't exist",
  "code": "4",
  "target": "uuid"
}
}
```

---

### Retrieve details for all subnets

GET `/network/ip/subnets`

**Introduced In:** 9.11

Retrieves details for all subnets.

#### Related ONTAP Commands

- `network subnet show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| available_count | integer | query | False | Filter by available_count |
| subnet.family | string | query | False | Filter by subnet.family |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| subnet.address | string | query | False | Filter by subnet.address |
| subnet.netmask | string | query | False | Filter by subnet.netmask |
| name | string | query | False | Filter by name |
| total_count | integer | query | False | Filter by total_count |
| broadcast_domain.uuid | string | query | False | Filter by broadcast_domain.uuid |
| broadcast_domain.name | string | query | False | Filter by broadcast_domain.name |
| ip_ranges.family | string | query | False | Filter by ip_ranges.family |
| ip_ranges.start | string | query | False | Filter by ip_ranges.start |
| ip_ranges.end | string | query | False | Filter by ip_ranges.end |
| gateway | string | query | False | Filter by gateway |
| used_count | integer | query | False | Filter by used_count |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| uuid | string | query | False | Filter by uuid |
| available_ip_ranges.family | string | query | False | Filter by available_ip_ranges.family |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| available_ip_ranges.start | string | query | False | Filter by available_ip_ranges.start |
| available_ip_ranges.end | string | query | False | Filter by available_ip_ranges.end |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[ip_subnet] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "available_count": 0,
      "available_ip_ranges": [
        {
          "end": "10.10.10.7",
          "family": "string",
          "start": "10.10.10.7"
        }
      ],
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "fail_if_lifs_conflict": true,
      "gateway": "10.1.1.1",
      "ip_ranges": [
        {
          "end": "10.10.10.7",
          "family": "string",
          "start": "10.10.10.7"
        }
      ],
      "ipspace": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "subnet1",
      "subnet": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "total_count": 0,
      "used_count": 0,
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_address_range

IP address range

| Name | Type | Description |
|---|---|---|
| end | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| start | string | IPv4 or IPv6 address |

broadcast_domain

The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

ipspace

The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_subnet

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| broadcast_domain | broadcast_domain | The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |
| ipspace | ipspace | The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| name | string | Subnet name |

| Name | Type | Description |
|---|---|---|
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new named subnet**

POST `/network/ip/subnets`

**Introduced In:** 9.11

Creates a new named subnet.

**Required properties**

- `name` - Name of the subnet to create.
- `broadcast_domain` - Broadcast domain containing the subnet.
- `ipspace` - IPspace containing the subnet. Required only if `broadcast_domain.uuid` is not provided.
- `subnet.address` - IP address for the subnet.
- `subnet.netmask` - IP netmask of the subnet.

#### Recommended property values

#### Default property values

If not specified in POST, the following default property values are assigned:

- `gateway` - no gateway
- `ip_ranges` - empty
- `fail_if_lifs_conflict` - *true*

#### Related ONTAP commands

- `network subnet create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| broadcast_domain | broadcast_domain | The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| fail_if_lifs_conflict | boolean | This action will fail if any existing interface is using an IP address in the ranges provided. Set this to false to associate any manually addressed interfaces with the subnet and allow the action to succeed. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| name | string | Subnet name |
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

**Example request**

```
{
  "available_count": 0,
  "available_ip_ranges": [
    {
      "end": "10.10.10.7",
      "start": "10.10.10.7"
    }
  ],
  "broadcast_domain": {
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "gateway": "10.1.1.1",
  "ip_ranges": [
    {
      "end": "10.10.10.7",
      "start": "10.10.10.7"
    }
  ],
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "subnet1",
  "subnet": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "total_count": 0,
  "used_count": 0,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[ip_subnet] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "available_count": 0,
      "available_ip_ranges": [
        {
          "end": "10.10.10.7",
          "start": "10.10.10.7"
        }
      ],
      "broadcast_domain": {
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "gateway": "10.1.1.1",
      "ip_ranges": [
        {
          "end": "10.10.10.7",
          "start": "10.10.10.7"
        }
      ],
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "subnet1",
      "subnet": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "total_count": 0,
      "used_count": 0,
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1377658 | Invalid gateway for subnet in IPspace. |
| 1377659 | Subnet would overlap with existing subnet named in IPspace. |
| 1377660 | A subnet with the name already exists in the IPspace. |
| 1377661 | Subnet in IPspace cannot use subnet address because that address is already used by another subnet in the same IPspace. |
| 1377662 | The IP range address is not within the subnet in IPspace. |
| 1377663 | The specified IP address range of subnet in IPspace contains an address already in use by a LIF. |
| 1377664 | The specified IP address range of subnet in IPspace contains an address already in use by the Service Processor. |
| 1377673 | The addresses provided must have the same address family. |
| 1377675 | The netmask of the interface did not match the netmask of the subnet. |
| 1377681 | Cannot update LIF associations for LIF. The broadcast domain of the LIF does not match the broadcast domain of the subnet. |
| 1377682 | IPv6 is not enabled in the cluster. |
| 1966269 | IPv4 Addresses must have a prefix length between 1 and 32. |
| 1967082 | The specified ipspace.name does not match the IPspace name of specified ipspace.uuid |
| 53282568 | The subnet.address must be specified together with subnet.netmask. |
| 53282569 | The specified subnet.netmask is not valid. |
| 53282570 | Each pair of ranges must have ip_ranges.start less than or equal to ip_ranges.end. |

| Error Code | Description |
| --- | --- |
| 53282571 | The ip_ranges.start and ip_ranges.end fields must have the same number of items. |
| 53282573 | Broadcast domain is a required parameter. The broadcast_domain.name and/or the broadcast_domain.uuid must be specified. |
| 53282574 | The specified broadcast_domain and ipspace parameters do not match. |
| 53282575 | Operation might have left configuration in an inconsistent state. Unable to set UUID for created entry. |
| 53282576 | The specified ipspace.uuid is invalid. |
| 53282577 | The specified broadcast_domain.uuid is invalid. |
| 53282578 | The specified broadcast_domain.name does not match the IPspace name of specified broadcast_domain.uuid |
| 53282579 | Missing the ipspace.name or ipspace.uuid parameter. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip_address_range

IP address range

| Name | Type | Description |
|---|---|---|
| end | string | IPv4 or IPv6 address |
| start | string | IPv4 or IPv6 address |

broadcast_domain

The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

ipspace

The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_subnet

| Name | Type | Description |
| --- | --- | --- |
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| broadcast_domain | broadcast_domain | The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| fail_if_lifs_conflict | boolean | This action will fail if any existing interface is using an IP address in the ranges provided. Set this to false to associate any manually addressed interfaces with the subnet and allow the action to succeed. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |
| ipspace | ipspace | The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| name | string | Subnet name |
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an IP subnet

DELETE `/network/ip/subnets/{uuid}`

**Introduced In:** 9.11

Deletes an IP subnet.

#### Related ONTAP commands

- `network subnet delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP subnet UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

Fill error codes below.
ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1377663 | The specified IP address range of subnet in IPspace contains an address already in use by a LIF. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve IP subnet details

GET `/network/ip/subnets/{uuid}`

**Introduced In:** 9.11

Retrieves details for a specific IP subnet.

#### Related ONTAP commands

- `network subnet show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP subnet UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| broadcast_domain | broadcast_domain | The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |
| ipspace | ipspace | The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST. |
| name | string | Subnet name |
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "available_count": 0,
  "available_ip_ranges": [
    {
      "end": "10.10.10.7",
      "family": "string",
      "start": "10.10.10.7"
    }
  ],
  "broadcast_domain": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "fail_if_lifs_conflict": true,
  "gateway": "10.1.1.1",
  "ip_ranges": [
    {
      "end": "10.10.10.7",
      "family": "string",
      "start": "10.10.10.7"
    }
  ],
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "subnet1",
  "subnet": {
    "address": "10.10.10.7",
```

```
    "family": "string",
    "netmask": "24"
  },
  "total_count": 0,
  "used_count": 0,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ip_address_range

IP address range

| Name | Type | Description |
|---|---|---|
| end | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| start | string | IPv4 or IPv6 address |

broadcast_domain

The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

ipspace

The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an IP subnet

PATCH `/network/ip/subnets/{uuid}`

**Introduced In:** 9.11

Updates an IP subnet.

#### Related ONTAP commands

- `network subnet modify`

- `network subnet rename`
- `network subnet add-ranges`
- `network subnet remove-ranges`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IP subnet UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| fail_if_lifs_conflict | boolean | This action will fail if any existing interface is using an IP address in the ranges provided. Set this to false to associate any manually addressed interfaces with the subnet and allow the action to succeed. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |
| name | string | Subnet name |
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

#### Example request

```
{
  "available_count": 0,
  "available_ip_ranges": [
    {
      "end": "10.10.10.7",
      "start": "10.10.10.7"
    }
  ],
  "gateway": "10.1.1.1",
  "ip_ranges": [
    {
      "end": "10.10.10.7",
      "start": "10.10.10.7"
    }
  ],
  "name": "subnet1",
  "subnet": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "total_count": 0,
  "used_count": 0,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

Fill error codes below.
ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1377658 | Invalid gateway for subnet in IPspace. |
| 1377659 | Subnet would overlap with existing subnet named in IPspace. |

| Error Code | Description |
| --- | --- |
| 1377660 | A subnet with the name already exists in the IPspace. |
| 1377661 | Subnet in IPspace cannot use subnet address because that address is already used by subnet in the same IPspace. |
| 1377662 | The IP range address is not within the subnet in IPspace. |
| 1377663 | The specified IP address range of subnet in IPspace contains an address already in use by a LIF. |
| 1377664 | The specified IP address range of subnet in IPspace contains an address already in use by the Service Processor. |
| 1377669 | The specified gateway address exists in a configured range. |
| 1377673 | The addresses provided must have the same address family. |
| 1377674 | Modifying the netmask is not supported because the subnet has one or more interfaces associated with it. |
| 1377675 | The netmask of the interface did not match the netmask of the subnet. |
| 1377681 | Cannot update LIF associations for LIF. The broadcast domain of the LIF does not match the broadcast domain of the subnet. |
| 1966269 | IPv4 Addresses must have a prefix length between 1 and 32. |
| 1966778 | IPv6 addresses must have a prefix length between 1 and 127. |
| 53282568 | The subnet.address must be specified together with subnet.netmask. |
| 53282569 | The specified subnet.netmask is not valid. |
| 53282570 | Each pair of ranges must have ip_ranges.start less than or equal to ip_ranges.end. |
| 53282571 | The ip_ranges.start and ip_ranges.end fields must have the same number of items. |
| 53282572 | PATCH partially succeeded with error. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ip_address_range

IP address range

| Name | Type | Description |
|---|---|---|
| end | string | IPv4 or IPv6 address |
| start | string | IPv4 or IPv6 address |

broadcast_domain

The broadcast domain that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

ipspace

The IPspace that the subnet is associated with. Either the UUID or name must be supplied on POST.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_subnet

| Name | Type | Description |
| --- | --- | --- |
| available_count | integer | |
| available_ip_ranges | array[ip_address_range] | |
| fail_if_lifs_conflict | boolean | This action will fail if any existing interface is using an IP address in the ranges provided. Set this to false to associate any manually addressed interfaces with the subnet and allow the action to succeed. |
| gateway | string | The IP address of the gateway for this subnet. |
| ip_ranges | array[ip_address_range] | |
| name | string | Subnet name |
| subnet | ip_info | IP information |
| total_count | integer | |
| used_count | integer | |
| uuid | string | The UUID that uniquely identifies the subnet. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage network IPspaces

### Manage network IPspaces

#### Overview

An IPspace is an addressing domain within which each IP address is unique. The same address may appear in a different IPspace, but the matching addresses are considered to be distinct. SVMs and broadcast domains, and therefore IP interfaces and Ethernet ports, are associated with a single IPspace. This endpoint supports the following operations: GET (collection and instance), POST, PATCH, and DELETE.

#### Retrieving IPspace information

You can use the IPspaces GET API to retrieve all IPspaces configured in the cluster, including built-in and custom IPspaces, and specifically requested IPspaces.

#### Examples

### Retrieving a list of the IPspaces in the cluster

The following example returns the requested list of IPspaces configured in the cluster.

```
# The API:
/api/network/ipspaces

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ipspaces?fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "dcc7e79c-5acc-11e8-b9de-005056b42b32",
    "name": "Default",
    "_links": {
      "self": {
```

```
        "href": "/api/network/ipspaces/dcc7e79c-5acc-11e8-b9de-
005056b42b32"
      }
    }
  },
  {
    "uuid": "dfd3c1b2-5acc-11e8-b9de-005056b42b32",
    "name": "Cluster",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/dfd3c1b2-5acc-11e8-b9de-
005056b42b32"
      }
    }
  },
  {
    "uuid": "dedec1be-5aec-1eee-beee-0eee56be2b3e",
    "name": "Ipspace1",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/dedec1be-5aec-1eee-beee-
0eee56be2b3e"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/network/ipspaces?fields=*"
  }
}
}
```


#### Retrieving a specific IPspace in the cluster

The following example returns the specific IPspace requested. The system returns an error if there is no IPspace with the requested UUID.

```
# The API:
/api/network/ipspaces/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ipspaces/dfd3c1b2-5acc-11e8-
b9de-005056b42b32?fields=*" -H "accept: application/hal+json"

# The response:
{
  "uuid": "dcc7e79c-5acc-11e8-b9de-005056b42b32",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/dcc7e79c-5acc-11e8-b9de-005056b42b32"
    }
  }
}
```

#### Creating IPspaces

You can use the network IPspaces POST API to create IPspaces.

#### Example

### Creating an IPspace

The following output displays the record returned after the creation of an IPspace with the name "ipspace1".

```
# The API:
/api/network/ipspaces

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ipspaces?return_records=true"
-H  "accept: application/hal+json" -d "{  \"name\": \"ipspace2\"}"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "4165655e-0528-11e9-bd68-005056bb046a",
    "name": "ipspace2",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/4165655e-0528-11e9-bd68-
005056bb046a"
      }
    }
  }
]
}
```

#### Updating IPspaces

You can use the IPspaces PATCH API to update the attributes of the IPspace.

#### Example

### Updating the name of an IPspace

The following PATCH request is used to update the name of the IPspace from "ipspace2" to "ipspace20".

```
# The API:
/api/network/ipspaces/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ipspaces/4165655e-0528-11e9-
bd68-005056bb046a" -H  "accept: application/hal+json" -d "{  \"name\":
\"ipspace20\"}"
```

#### Deleting IPspaces

You can use the IPspaces DELETE API to delete an IPspace.

#### Example

### Deleting an IPspace

The following DELETE request is used to delete an IPspace.

```
# The API:
/api/network/ipspaces/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ipspaces/4165655e-0528-11e9-
bd68-005056bb046a" -H  "accept: application/hal+json" -H  "Content-Type:
application/json"
```

## Retrieve IPspaces for a cluster

GET `/network/ipspaces`

**Introduced In:** 9.6

Retrieves a collection of IPspaces for the entire cluster.

#### Related ONTAP commands

- `network ipspace show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[ipspace] | |

**Example response**

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "ipspace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create a new domain with unique IP addresses

POST `/network/ipspaces`

**Introduced In:** 9.6

Creates a new domain within which IP addresses are unique. SVMs, ports, and networks are scoped to a single IPspace.

#### Required properties

- `name` - Name of the IPspace to create.

#### Related ONTAP commands

- `network ipspace create`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

#### Example request

```
{
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1377261 | An IPspace with the same name already exists. |
| 1966449 | Invalid IPspace name in MetroCluster configurations. |
| 1966586 | The specified IPspace name is invalid because it is already used by a peered SVM. |
| 1967102 | A POST operation might have left the configuration in an inconsistent state. Check the configuration. |
| 9240587 | Name cannot be empty. |
| 9240588 | The name is too long. |
| 9240589 | Invalid character in name. |
| 9240590 | The name is reserved by the system. |
| 9240591 | The name is not valid. The name is already in use by a cluster node, Vserver, or it is the name of the local cluster. |
| 53281576 | Cannot create IPspace because the maximum number of custom IPspaces has already been reached on the cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an IPspace object

DELETE `/network/ipspaces/{uuid}`

**Introduced In:** 9.6

Deletes an IPspace object.

#### Related ONTAP commands

- `network ipspace delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IPspace UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966309 | The IPspace cannot be removed because it is being used by the audit log forwarding feature. |
| 1966333 | IPspace can't be removed because it has SVMs assigned. |
| 1969180 | Cannot delete an IPspace with an existing broadcast domain. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve information about an IPspace

GET `/network/ipspaces/{uuid}`

**Introduced In:** 9.6

Retrieves information about a specific IPspace.

#### Related ONTAP commands

- `network ipspace show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IPspace UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an IPspace object

PATCH `/network/ipspaces/{uuid}`

**Introduced In:** 9.6

Updates an IPspace object.

#### Related ONTAP commands

- `network ipspace rename`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | IPspace UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

### Example request

```
{
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1377256 | Another IPspace with the same name exists. |
| 1377259 | System-defined IPspaces cannot be modified or removed. |
| 1966449 | Invalid IPspace name in MetroCluster configurations. |
| 9240587 | IPspace name cannot be empty. |
| 9240588 | The specified name is too long. |
| 9240589 | Invalid character in specified name. |
| 9240590 | Invalid name, the name is reserved by the system. |

| Error Code | Description |
| --- | --- |
| 9240591 | Invalid name. The name is already in use by a cluster node, Vserver, or it is the name of the local cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

ipspace

| Name | Type | Description |
| --- | --- | --- |
| name | string | IPspace name |
| uuid | string | The UUID that uniquely identifies the IPspace. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Object-store

## Manage S3 buckets

### Manage S3 buckets

#### Overview

An S3 bucket is a container of objects. Each bucket defines an object namespace. S3 server requests specify objects using a bucket-name and object-name pair. An object consists of data, along with optional metadata and access controls, that is accessible using a name. An object resides within a bucket. There can be more than one bucket in an S3 server. Buckets that are created for the server are associated with an S3 user that is created on the S3 server.

#### Examples

### Retrieving all fields for all S3 buckets of a cluster

```
# The API:
/api/protocols/s3/buckets

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/buckets?fields=**&return_records=true" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
    "name": "bucket-2",
    "volume": {
      "name": "fg_oss_1558514455",
      "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
    },
    "size": 107374182400,
    "logical_used_size": 157286400,
    "encryption": {
      "enabled": false
    },
    "comment": "S3 bucket.",
    "qos_policy": {
      "min_throughput_iops": 0,
```

```
      "min_throughput_mbps": 0,
      "max_throughput_iops": 1000,
      "max_throughput_mbps": 0,
      "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
      "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
    },
    "snapshot-policy": {
      "name": "default-1weekly",
      "uuid": "f9c5f090-4ac8-11ef-ba24-005056a7ceb6"
    },
  },
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
    "name": "bucket-1",
    "volume": {
      "name": "fg_oss_1558690256",
      "uuid": "a36a1ea7-7e06-11e9-97e8-0050568ea123"
    },
    "size": 107374182400,
    "logical_used_size": 0,
    "encryption": {
      "enabled": false
    },
    "comment": "bucket2",
    "qos_policy": {
      "min_throughput_iops": 0,
      "min_throughput_mbps": 0,
      "max_throughput_iops": 1000,
      "max_throughput_mbps": 0,
      "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
      "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
    }
  },
  {
    "svm": {
      "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
      "name": "vs2"
    },
    "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
    "name": "bucket-3",
    "volume": {
      "name": "fg_oss_1558690257",
```

```
      "uuid": "a46a1ea7-7e06-11e9-97e8-0050568ea123"
    },
    "size": 107374182400,
    "logical_used_size": 1075838976,
    "encryption": {
      "enabled": false
    },
    "comment": "bucket3",
    "qos_policy": {
      "min_throughput_iops": 0,
      "min_throughput_mbps": 0,
      "max_throughput_iops": 1000,
      "max_throughput_mbps": 0,
      "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
      "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
    },
    "policy": {
      "statements": [
        {
          "effect": "allow",
          "actions": [
            "*"
          ],
          "principals": [
            "Alice"
          ],
          "resources": [
            "bucket-3",
            "bucket-3/*"
          ],
          "sid": "fullAccessForAliceToBucket",
          "conditions": [
            {
              "operator": "ip_address",
              "source_ips": [
                "1.1.1.1/10"
              ]
            }
          ]
        }
      ]
    },
    "cors": {
      "rules": [
        {
          "id": "string",
```

```
          "allowed_origins": [
            "http://www.example.com"
          ],
          "allowed_methods": [
            "PUT",
            "DELETE"
          ],
          "allowed_headers": [
            "x-amz-request-id"
          ],
          "expose_headers": [
            "http://www.example.com"
          ],
          "max_age_seconds": 1024
        }
      ]
    }
  }
],
"num_records": 3
}
```


**Retrieving all S3 buckets of a cluster ordered by size**

```
# The API:
/api/protocols/s3/buckets

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/buckets?return_records=true&order_by=size" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "754389d0-7e13-11e9-bfdc-0050568ea123",
    "name": "bb1",
    "size": 107374182400
  },
  {
    "svm": {
      "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
      "name": "vs2"
    },
    "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
    "name": "bb2",
    "size": 107374182400
  },
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
    "name": "bucket-1",
    "size": 107374182400
  }
],
"num_records": 3
}
```

**Retrieving all S3 buckets of a cluster with name "bb2"**

```
# The API:
/api/protocols/s3/buckets

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/buckets?name=bb2&return_records=true" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "087d940e-7e15-11e9-bfdc-0050568ea123",
    "name": "bb2"
  },
  {
    "svm": {
      "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
      "name": "vs2"
    },
    "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
    "name": "bb2"
  }
],
"num_records": 2
}
```

**Retrieving the specified bucket associated with an SVM**

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/buckets/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/527812ab-7c6d-11e9-97e8-0050568ea123" -H "accept:
application/json"

# The response:
{
"svm": {
  "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
  "name": "vs1"
},
"uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
"name": "bucket-2",
"volume": {
  "name": "fg_oss_1558514455",
  "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
},
"size": 107374182400,
"logical_used_size": 157286400,
"encryption": {
  "enabled": false
},
"comment": "S3 bucket.",
"qos_policy": {
  "min_throughput_iops": 0,
  "min_throughput_mbps": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
}
}
```

**Creating an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
\"constituents_per_aggregate\": 4, \"name\": \"bucket-3\", \"svm\": {
\"name\": \"vs1\" } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Creating an S3 bucket along with QoS policy for an SVM**

```
# The API:
/api/protocols/s3/buckets

# The call:
curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"comment\": \"S3 bucket.\", \"name\": \"bucket-3\", \"svm\": { \"name\":
\"vs1\" },  \"qos_policy\": { \"min_throughput_iops\": 0,
\"min_throughput_mbps\": 0, \"max_throughput_iops\": 1000000,
\"max_throughput_mbps\": 900000, \"uuid\": \"02d07a93-6177-11ea-b241-
000c293feac8\", \"name\":
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Creating an S3 bucket along with policies and conditions for an SVM**

```
# The API:
/api/protocols/s3/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
\"constituents_per_aggregate\": 4, \"name\": \"bucket-3\", \"policy\": {
\"statements\": [ { \"actions\": [ \"GetObject\" ], \"conditions\": [ {
\"operator\": \"ip_address\", \"source_ips\": [ \"1.1.1.1/23\",
\"1.2.2.2/20\" ] } ], \"effect\": \"allow\", \"resources\": [ \"bucket-
3/policies/examples/*\" ], \"sid\": \"AccessToGetObjectForAllUsersofSVM\"
}, { \"actions\": [ \"*Object\" ], \"effect\": \"deny\", \"principals\": [
\"mike\" ], \"resources\": [ \"bucket-3/policy-docs/*\", \"bucket-
3/confidential-*\" ], \"sid\": \"DenyAccessToObjectForMike\" }, {
\"actions\": [ \"GetObject\" ], \"effect\": \"allow\", \"principals\": [
\"*\" ], \"resources\": [ \"bucket-3/readme\" ], \"sid\":
\"AnonymousAccessToGetObjectForUsers\" } ] }, \"svm\": { \"uuid\":
\"259b4e46-2d33-11ea-9145-005056bbbec1\" } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
```

```
    "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
      }
    }
  }
  }
```


### Creating an S3 bucket and CORS rules for an SVM

```
# The API:
/api/protocols/s3/buckets

# The call:

curl -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
\"constituents_per_aggregate\": 4, \"name\": \"bucket-3\", \"cors\": {
\"rules\": [{ \"allowed_headers\": [ \"x-amz-request-id\" ],
\"allowed_methods\": [ \"PUT\", \"DELETE\" ], \"allowed_origins\": [
\"http://www.example.com\" ], \"expose_headers\": [
\"http://www.example.com\" ], \"id\": \"id1\", \"max_age_seconds\": 1024
}]}, \"svm\": { \"uuid\": \"259b4e46-2d33-11ea-9145-005056bbbec1\" }}"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Creating an S3 bucket along with a bucket lifecycle management rule for an SVM**

```
# The API:
/api/protocols/s3/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
\"constituents_per_aggregate\": 4, \"name\": \"bucket-4\",
\"lifecycle_management\": { \"rules\": [ { \"name\": \"rule1\",
\"expiration\": { \"object_age_days\" : \"1000\" },
\"abort_incomplete_multipart_upload\" : { \"after_initiation_days\" : 200
}, \"object_filter\": { \"prefix\" : \"obj1*/\" ,  \"size_greater_than\" :
\"1000\" } }, { \"name\": \"rule2\", \"object_filter\": {
\"size_greater_than\" : \"50\" }, \"expiration\": { \"object_age_days\" :
\"5000\" } } ] } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 18 April 2022 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-4
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-4",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
```

```
    }
  }
}
}
```

#### Creating an S3 bucket with a snapshot-policy

```
# The API:
/api/protocols/s3/buckets

# The call:
curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"comment\":
\"S3 snapshot policy bucket.\", \"snapshot_policy\": { \"name\":
\"default-1weekly\", \"uuid\": \"f9c5f090-4ac8-11ef-ba24-005056a7ceb6\" },
\"name\": \"bucket-7\", \"svm.uuid\": \"8c38f10b-4871-11ef-aab5-
005056a7ceb6\" }"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 25 Jul 2024 17:16:15 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/8c38f10b-4871-11ef-aab5-
005056a7ceb6/buckets/9f21f404-4aa4-11ef-ba24-005056a7ceb6
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "99978136-4aa9-11ef-ba24-005056a7ceb6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/99978136-4aa9-11ef-ba24-005056a7ceb6"
    }
  }
}
}
```

#### Updating an S3 bucket for an SVM

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/259b4e46-2d33-
11ea-9145-005056bbbec1/376a2efd-2d4d-11ea-9c30-
005056bb883a?return_records=true" -H "accept:
application/json?return_records=true" -H "Content-Type: application/json"
-d "{ \"comment\": \"Bucket modified.\", \"size\": 111111111111,
\"qos_policy\": { \"min_throughput_iops\": 0, \"min_throughput_mbps\": 0,
\"max_throughput_iops\": 1000000, \"max_throughput_mbps\": 900000,
\"uuid\": \"02d07a93-6177-11ea-b241-000c293feac8\", \"name\":
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" }}"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating an S3 bucket policy for an SVM**

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/259b4e46-2d33-
11ea-9145-005056bbbec1/376a2efd-2d4d-11ea-9c30-
005056bb883a?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ {
\"actions\": [ \"*\" ], \"conditions\": [ { \"operator\": \"ip_address\",
\"source_ips\": [ \"1.1.1.5/23\" ] } ], \"effect\": \"allow\",
\"resources\": [ \"*\" ], \"sid\": \"fullAccessForAllPrincipalsToBucket\"}
] } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating an S3 bucket CORS configuration for an SVM**

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/259b4e46-2d33-
11ea-9145-005056bbbec1/376a2efd-2d4d-11ea-9c30-
005056bb883a?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"cors\": { \"rules\": [{
\"allowed_headers\": [ \"x-amz-request-id\" ], \"allowed_methods\": [
\"PUT\", \"DELETE\" ], \"allowed_origins\": [ \"http://www.example.com\"
], \"expose_headers\": [ \"http://www.example.com\" ], \"id\": \"id1\",
\"max_age_seconds\": 1024 } ] } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating the snapshot-policy for an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/8c38f10b-4871-
11ef-aab5-005056a7ceb6/eec0d90e-4934-11ef-947d-005056a7ceb6" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"snapshot_policy\": { \"name\": \"default-1weekly\", \"uuid\":
\"f9c5f090-4ac8-11ef-ba24-005056a7ceb6\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 25 Jul 2024 17:26:17 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "003a4cad-4aab-11ef-ba24-005056a7ceb6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/003a4cad-4aab-11ef-ba24-005056a7ceb6"
    }
  }
}
}
```

**Deleting an S3 bucket for a specified SVM**

```
# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -iku admin:<password> -X DELETE "https://<mgmt-
ip>/api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/98528221-2d52-11ea-892e-005056bbbec1?return_records=true" -H
"accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:40:17 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "b3af4a54-7e18-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b3af4a54-7e18-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

### Retrieve all S3 buckets for all SVMs

GET `/protocols/s3/buckets`

**Introduced In:** 9.7

Retrieves all S3 buckets for all SVMs. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket show`
- `vserver object-store-server bucket policy statement show`
- `vserver object-store-server bucket policy-statement-condition show`
- `vserver object-store-server bucket lifecycle-management-rule show`
- `vserver object-store-server bucket cors-rule show`

**Learn more**

- DOC /protocols/s3/buckets

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| is_nas_path_mutable | boolean | query | False | Filter by is_nas_path_mutable<br>• Introduced in: 9.17 |
| nas_path | string | query | False | Filter by nas_path<br>• Introduced in: 9.12 |
| encryption.enabled | boolean | query | False | Filter by encryption.enabled |
| allowed | boolean | query | False | Filter by allowed<br>• Introduced in: 9.12 |
| protection_status.destination.is_ontap | boolean | query | False | Filter by protection_status.destination.is_ontap<br>• Introduced in: 9.10 |
| protection_status.destination.is_cloud | boolean | query | False | Filter by protection_status.destination.is_cloud<br>• Introduced in: 9.10 |
| protection_status.destination.is_external_cloud | boolean | query | False | Filter by protection_status.destination.is_external_cloud<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protection_status.is_protected | boolean | query | False | Filter by protection_status.is_protected<br>• Introduced in: 9.10 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| retention.default_period | string | query | False | Filter by retention.default_period<br>• Introduced in: 9.14 |
| retention.mode | string | query | False | Filter by retention.mode<br>• Introduced in: 9.14 |
| role | string | query | False | Filter by role<br>• Introduced in: 9.10 |
| size | integer | query | False | Filter by size<br>• Max value: 67553994410557440<br>• Min value: 107374182400 |
| logical_used_size | integer | query | False | Filter by logical_used_size |
| uuid | string | query | False | Filter by uuid |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| qos_policy.min_throughput_mbps | integer | query | False | Filter by qos_policy.min_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| qos_policy.max_throughput_iops | integer | query | False | Filter by qos_policy.max_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| qos_policy.max_throughput_mbps | integer | query | False | Filter by qos_policy.max_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| qos_policy.name | string | query | False | Filter by qos_policy.name<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| qos_policy.min_throughput_iops | integer | query | False | Filter by qos_policy.min_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| qos_policy.uuid | string | query | False | Filter by qos_policy.uuid<br>• Introduced in: 9.8 |
| qos_policy.max_throughput | string | query | False | Filter by qos_policy.max_throughput<br>• Introduced in: 9.17 |
| qos_policy.min_throughput | string | query | False | Filter by qos_policy.min_throughput<br>• Introduced in: 9.17 |
| versioning_state | string | query | False | Filter by versioning_state<br>• Introduced in: 9.11 |
| is_consistent_etag | boolean | query | False | Filter by is_consistent_etag<br>• Introduced in: 9.17 |
| audit_event_selector.permission | string | query | False | Filter by audit_event_selector.permission<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| audit_event_selector.access | string | query | False | Filter by audit_event_selector.access<br>• Introduced in: 9.10 |
| type | string | query | False | Filter by type<br>• Introduced in: 9.12 |
| lifecycle_management.rules.bucket_name | string | query | False | Filter by lifecycle_management.rules.bucket_name<br>• Introduced in: 9.14<br>• maxLength: 63<br>• minLength: 3 |
| lifecycle_management.rules.non_current_version_expiration.non_current_days | integer | query | False | Filter by lifecycle_management.rules.non_current_version_expiration.non_current_days<br>• Introduced in: 9.13 |
| lifecycle_management.rules.non_current_version_expiration.new_non_current_versions | integer | query | False | Filter by lifecycle_management.rules.non_current_version_expiration.new_non_current_versions<br>• Introduced in: 9.13 |
| lifecycle_management.rules.uuid | string | query | False | Filter by lifecycle_management.rules.uuid<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lifecycle_management.rules.expiration.expired_object_delete_marker | boolean | query | False | Filter by lifecycle_management.rules.expiration.expired_object_delete_marker<br>• Introduced in: 9.13 |
| lifecycle_management.rules.expiration.object_expiry_date | string | query | False | Filter by lifecycle_management.rules.expiration.object_expiry_date<br>• Introduced in: 9.13 |
| lifecycle_management.rules.expiration.object_age_days | integer | query | False | Filter by lifecycle_management.rules.expiration.object_age_days<br>• Introduced in: 9.13 |
| lifecycle_management.rules.name | string | query | False | Filter by lifecycle_management.rules.name<br>• Introduced in: 9.13<br>• maxLength: 256<br>• minLength: 0 |
| lifecycle_management.rules.svm.name | string | query | False | Filter by lifecycle_management.rules.svm.name<br>• Introduced in: 9.14 |
| lifecycle_management.rules.svm.uuid | string | query | False | Filter by lifecycle_management.rules.svm.uuid<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lifecycle_management.rules.object_filter.prefix | string | query | False | Filter by lifecycle_management.rules.object_filter.prefix<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.object_filter.size_greater_than | integer | query | False | Filter by lifecycle_management.rules.object_filter.size_greater_than<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.object_filter.size_less_than | integer | query | False | Filter by lifecycle_management.rules.object_filter.size_less_than<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.object_filter.tags | string | query | False | Filter by lifecycle_management.rules.object_filter.tags<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.enabled | boolean | query | False | Filter by lifecycle_management.rules.enabled<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days | integer | query | False | Filter by lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.statements.principals | string | query | False | Filter by policy.statements.principals<br>• Introduced in: 9.8 |
| policy.statements.sid | string | query | False | Filter by policy.statements.sid<br>• Introduced in: 9.8<br>• maxLength: 256<br>• minLength: 0 |
| policy.statements.resources | string | query | False | Filter by policy.statements.resources<br>• Introduced in: 9.8 |
| policy.statements.effect | string | query | False | Filter by policy.statements.effect<br>• Introduced in: 9.8 |
| policy.statements.actions | string | query | False | Filter by policy.statements.actions<br>• Introduced in: 9.8 |
| policy.statements.conditions.operator | string | query | False | Filter by policy.statements.conditions.operator<br>• Introduced in: 9.8 |
| policy.statements.conditions.delimiters | string | query | False | Filter by policy.statements.conditions.delimiters<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.statements.conditions.prefixes | string | query | False | Filter by policy.statements.conditions.prefixes<br>• Introduced in: 9.8 |
| policy.statements.conditions.usernames | string | query | False | Filter by policy.statements.conditions.usernames<br>• Introduced in: 9.8 |
| policy.statements.conditions.source_ips | string | query | False | Filter by policy.statements.conditions.source_ips<br>• Introduced in: 9.8 |
| policy.statements.conditions.max_keys | integer | query | False | Filter by policy.statements.conditions.max_keys<br>• Introduced in: 9.8 |
| name | string | query | False | Filter by name<br>• maxLength: 63<br>• minLength: 3 |
| snapshot_restore.objects_remaining | integer | query | False | Filter by snapshot_restore.objects_remaining<br>• Introduced in: 9.18 |
| snapshot_restore.state | string | query | False | Filter by snapshot_restore.state<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| snapshot_restore.snapshot | string | query | False | Filter by snapshot_restore.snapshot<br>• Introduced in: 9.18 |
| snapshot_restore.progress | integer | query | False | Filter by snapshot_restore.progress<br>• Introduced in: 9.18 |
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid<br>• Introduced in: 9.16 |
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name<br>• Introduced in: 9.16 |
| cors.rules.allowed_headers | string | query | False | Filter by cors.rules.allowed_headers<br>• Introduced in: 9.16 |
| cors.rules.id | string | query | False | Filter by cors.rules.id<br>• Introduced in: 9.16<br>• maxLength: 256<br>• minLength: 0 |
| cors.rules.max_age_seconds | integer | query | False | Filter by cors.rules.max_age_seconds<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| cors.rules.allowed_origins | string | query | False | Filter by cors.rules.allowed_origins<br><br>• Introduced in: 9.16 |
| cors.rules.expose_headers | string | query | False | Filter by cors.rules.expose_headers<br><br>• Introduced in: 9.16 |
| cors.rules.allowed_methods | string | query | False | Filter by cors.rules.allowed_methods<br><br>• Introduced in: 9.16 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_bucket] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "audit_event_selector": {
        "access": "string",
        "permission": "string"
      },
      "comment": "S3 bucket.",
      "cors": {
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "allowed_headers": [
              "x-amz-request-id"
            ],
            "allowed_methods": [
              "PUT",
              "DELETE"
            ],
            "allowed_origins": [
              "http://www.example.com"
            ],
            "expose_headers": [
              "x-amz-date"
            ],
            "id": "string",
            "max_age_seconds": 1024
          }
        ]
      },
      "lifecycle_management": {
```

```
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "abort_incomplete_multipart_upload": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              }
            },
            "bucket_name": "bucket1",
            "expiration": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "object_age_days": 100,
              "object_expiry_date": "2039-09-23 00:00:00 +0000"
            },
            "name": "string",
            "non_current_version_expiration": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              }
            },
            "object_filter": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "prefix": "/logs",
              "size_greater_than": 10240,
              "size_less_than": 10485760,
              "tags": [
                "project1=projA",
                "project2=projB"
              ]
            },
```

```
            "svm": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
          }
        ]
      },
      "logical_used_size": 0,
      "name": "bucket1",
      "nas_path": "/",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  1000
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ],
```

```
            "effect": "allow",
            "principals": [
              "user1",
              "group/grp1",
              "nasgroup/group1"
            ],
            "resources": [
              "bucket1",
              "bucket1/*"
            ],
            "sid": "Full_Access_To_User1!"
          }
        ]
      },
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "retention": {
```

```
        "default_period": "P10Y",
        "mode": "governance"
      },
      "role": "string",
      "size": 214748364800,
      "snapshot_policy": {
        "name": "default-1weekly",
        "uuid": "3675af31-431c-12fa-114a-20675afebc12"
      },
      "snapshot_restore": {
        "snapshot": "string",
        "state": "string"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "s3",
      "use_mirrored_aggregates": true,
      "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
      "versioning_state": "enabled",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |

| Name | Type | Description |
| --- | --- | --- |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value:<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value:<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create the S3 bucket configuration for an SVM

POST `/protocols/s3/buckets`

**Introduced In:** 9.8

Creates the S3 bucket configuration of an SVM.

#### Important notes

- Each SVM can have one or more bucket configurations.
- Aggregate lists should be specified explicitly. If not specified, then the bucket is auto-provisioned as a FlexGroup volume.
- Constituents per aggregate specifies the number of components (or FlexVol volumes) per aggregate. Is specified only when an aggregate list is explicitly defined.
- An access policy can be created along with a bucket create. If creating an access policy fails, bucket configurations are saved and the access policy can be created using the PATCH endpoint.
- "qos_policy" can be specified if a bucket needs to be attached to a QoS group policy during creation time.
- "audit_event_selector" can be specified if a bucket needs to be specify access and permission type for auditing.
- A CORS configuration can be specified along with bucket creation.

#### Required properties

- `svm.uuid or svm.name` - Existing SVM in which to create the bucket configuration.
- `name` - Bucket name that is to be created.

#### Recommended optional properties

- `aggregates` - List of aggregates for the FlexGroup volume on which the bucket is hosted on.
- `constituents_per_aggregate` - Number of constituents per aggregate.
- `size` - Specifying the bucket size is recommended.
- `policy` - Specifying a policy enables users to perform operations on buckets; specifying the resource permissions is recommended.
- `qos_policy` - A QoS policy for buckets.
- `audit_event_selector` - Audit policy for buckets.
- `versioning_state` - Versioning state for buckets.
- `type` - Type of bucket.
- `nas_path` - NAS path to which the bucket corresponds to.

- `use_mirrored_aggregates` - Specifies whether mirrored aggregates are selected when provisioning a FlexGroup volume.
- `lifecycle_management` - Object store server lifecycle management policy.
- `retention.mode` - Object lock mode supported on the bucket.
- `retention.default_period` - Specifies the duration of default-retention applicable for objects on the object store bucket.
- `cors` - Specifying CORS rules enables the bucket to service the cross-origin requests.
- `snapshot_policy` - Snapshot policy for the bucket.
- `is_nas_path_mutable` - Specifies whether the NAS bucket mapping with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount.

#### Default property values

- `size` - 800MB
- `comment` - ""
- `aggregates` - No default value.
- `constituents_per_aggregate` - *4* , if an aggregates list is specified. Otherwise, no default value.
- `policy.statements.actions` - GetObject, PutObject, DeleteObject, ListBucket, ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging, PutObjectTagging, DeleteObjectTagging, GetBucketVersioning, PutBucketVersioning.
- `policy.statements.principals` - all S3 users and groups in the SVM or the NAS groups.
- `policy.statements.resources` - all objects in the bucket.
- `policy.statements.conditions` - list of bucket policy conditions.
- `versioning_state` - disabled.
- `use_mirrored_aggregates` - *true* for a MetroCluster configuration and *false* for a non-MetroCluster configuration.
- `type` - S3
- `retention.mode` - no_lock

#### Related ONTAP commands

- `vserver object-store-server bucket create`
- `vserver object-store-server bucket policy statement create`
- `vserver object-store-server bucket lifecycle-management-rule create`
- `vserver object-store-server bucket cors-rule create`

#### Learn more

- DOC /protocols/s3/buckets

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume. |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| constituents_per_aggregate | integer | Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |

| Name | Type | Description |
| --- | --- | --- |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| storage_service_level | string | Specifies the storage service level of the FlexGroup volume on which the bucket should be created. Valid values are "value", "performance" or "extreme". |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup. Only mirrored aggregates are used if this parameter is set to "true" and only unmirrored aggregates are used if this parameter is set to "false". The default value is "true" for a MetroCluster configuration and is "false" for a non-MetroCluster configuration. |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

**Example request**

```
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "constituents_per_aggregate": 4,
  "cors": {
    "rules": [
      {
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "lifecycle_management": {
    "rules": [
      {
        "bucket_name": "bucket1",
        "expiration": {
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "name": "string",
        "object_filter": {
```

```
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
    ]
  },
  "logical_used_size": 0,
  "name": "bucket1",
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
```

```
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
      }
    ]
  },
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y",
    "mode": "governance"
```

```
  },
  "role": "string",
  "size": 214748364800,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
    "state": "string"
  },
  "storage_service_level": "value",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "s3",
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
  "versioning_state": "enabled",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**


```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405777 + | "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason of failure}. "; |
| 92405785 + | "Bucket name "{bucket name}" contains invalid characters or invalid character combinations. Valid characters for a bucket name are 0-9, a-z, ".", and "-". Invalid character combinations are ".-", "-.", and "..". "; |
| 92405786 + | "Bucket name "{bucket name}" is not valid. Bucket names must have between 3 and 63 characters. "; |
| 92405811 + | "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |
| 92405812 + | "Failed to create the object store volume. Reason: {Reason for failure}."; |
| 92405819 + | "Cannot provision an object store server volume for bucket "{bucket name}" in SVM "{svm.name}" on the following aggregates because they are SnapLock aggregates: {List of aggregates.name}."; |
| 92405820 + | "Failed to check whether the aggregate "{aggregates.name}" is a FabricPool. Reason: {Reason for failure}."; |
| 92405821 + | "Cannot provision an object store server volume for bucket "{bucket name}" in SVM "{svm.name}" on the following aggregates because they are FabricPool: {List of aggregates.name}."; |
| 92405827 + | "Internal Error. Unable to generate object store volume name."; |

| Error code | Message |
| --- | --- |
| 92405857 + | "One or more aggregates must be specified if "constituents_per_aggregate" is specified."; |
| 92405858 + | "Failed to "create" the "bucket" because the operation is only supported on data SVMs."; |
| 92405859 + | "The specified "aggregates.uuid" "{aggregates.uuid}" does not exist."; |
| 92405860 + | "The specified "aggregates.name" "{aggregates.name}" and "aggregates.uuid" "{aggregates.uuid}" refer to different aggregates."; |
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92405863 + | "An error occurs when creating an access policy. The reason for failure is detailed in the error message."; |
| 92405863 + | "Failed to create lifecycle management rules for bucket "s3bucket1". Reason: {Reason of failure}. "; |
| 92405891 + | The resources specified in the access policy are not valid. Valid ways to specify a resource are *, <bucket-name>, <bucket-name>/…/…. Valid characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &, @,?, (, ), single quote, *, !, - and $. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |

| Name | Type | Description |
| --- | --- | --- |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
|---|---|---|
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume. |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| constituents_per_aggregate | integer | Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |

| Name | Type | Description |
| --- | --- | --- |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
|---|---|---|
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| storage_service_level | string | Specifies the storage service level of the FlexGroup volume on which the bucket should be created. Valid values are "value", "performance" or "extreme". |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup. Only mirrored aggregates are used if this parameter is set to "true" and only unmirrored aggregates are used if this parameter is set to "false". The default value is "true" for a MetroCluster configuration and is "false" for a non-MetroCluster configuration. |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the S3 bucket configuration for an SVM

DELETE `/protocols/s3/buckets/{svm.uuid}/{uuid}`

**Introduced In:** 9.8

Deletes the S3 bucket configuration of an SVM. An access policy is also deleted on an S3 bucket "delete" command.

#### Related ONTAP commands

- `vserver object-store-server bucket delete`
- `vserver object-store-server bucket policy statement delete`
- `vserver object-store-server bucket policy-statement-condition delete`
- `vserver object-store-server bucket lifecycle-management-rule delete`
- `vserver object-store-server bucket cors-rule delete`

#### Learn more

- DOC /protocols/s3/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405811 + | "Failed to delete bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |
| 92405858 + | "Failed to "delete" the "bucket" because the operation is only supported on data SVMs."; |
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92405779 + | "Failed to remove bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason for failure}. "; |
| 92405813 + | "Failed to delete the object store volume. Reason: {Reason for failure}."; |
| 92405864 + | "An error occurred when deleting an access policy. The reason for failure is detailed in the error message."; |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve the S3 bucket configuration for an SVM**

GET `/protocols/s3/buckets/{svm.uuid}/{uuid}`

**Introduced In:** 9.8

Retrieves the S3 bucket configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket show`
- `vserver object-store-server bucket policy statement show`
- `vserver object-store-server bucket policy-statement-condition show`
- `vserver object-store-server bucket lifecycle-management-rule show`
- `vserver object-store-server bucket cors-rule show`
- `vserver object-store-server bucket snapshot restore show`

#### Learn more

- DOC /protocols/s3/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |

| Name | Type | Description |
| --- | --- | --- |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |

| Name | Type | Description |
| --- | --- | --- |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

**Example response**


```
{
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "cors": {
    "rules": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "lifecycle_management": {
    "rules": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "abort_incomplete_multipart_upload": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
            }
          }
        },
        "bucket_name": "bucket1",
        "expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "name": "string",
        "non_current_version_expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        },
        "object_filter": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
```

```
    ]
  },
  "logical_used_size": 0,
  "name": "bucket1",
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
```

```
      }
    ]
  },
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y",
    "mode": "governance"
  },
  "role": "string",
  "size": 214748364800,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
```

```
    "state": "string"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "s3",
  "use_mirrored_aggregates": true,
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
  "versioning_state": "enabled",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |

| Name | Type | Description |
| --- | --- | --- |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| prefix | string | A prefix that is matched against object-names within a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the S3 bucket configuration for an SVM

PATCH `/protocols/s3/buckets/{svm.uuid}/{uuid}`

**Introduced In:** 9.8

Updates the S3 bucket configuration of an SVM.

#### Important notes

- The following fields can be modified for a bucket:
  - `comment` - Any information related to the bucket.
  - `size` - Bucket size.

- `policy` - An access policy for resources (buckets and objects) that defines their permissions. New policies are created after existing policies are deleted. To retain any of the existing policy statements, you need to specify those statements again. Also, policy conditions can be specified as part of a bucket policy.
- `qos_policy` - A QoS policy for buckets.
- `audit_event_selector` - Audit policy for buckets. None can be specified for both access and permission to remove an audit event selector.
- `versioning-state` - Versioning state of the buckets.
- `nas_path` - NAS path to which the bucket corresponds to.
- `retention.default_period` - Specifies the duration of default-retention applicable for objects on the object store bucket.
- `cors` - Specifying CORS rules enables the bucket to service the cross-origin requests. Note that the new CORS configuration specified will replace the existing one. If you need to retain any of the existing CORS rules, specify those rules again as part of the new CORS rules. To remove all the existing rules, specify an empty CORS configuration as input.
- `snapshot_policy` - Snapshot policy for the bucket.
- `restore_to` - Start a snapshot restore operation for S3 bucket

#### Related ONTAP commands

- `vserver object-store-server bucket modify`
- `vserver object-store-server bucket policy statement modify`
- `vserver object-store-server bucket policy-statement-condition modify`
- `vserver object-store-server bucket cors-rule create`
- `vserver object-store-server bucket cors-rule delete`
- `vserver object-store-server bucket snapshot restore start`

#### Learn more

- DOC /protocols/s3/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |

| Name | Type | Description |
| --- | --- | --- |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |

| Name | Type | Description |
| --- | --- | --- |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

**Example request**

```
{
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "cors": {
    "rules": [
      {
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "is_nas_path_mutable": true,
  "logical_used_size": 0,
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
```

```
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
      }
    ]
  },
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
```

```
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y"
  },
  "role": "string",
  "size": 214748364800,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
    "state": "string"
  },
  "type": "s3",
  "use_mirrored_aggregates": true,
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
  "versioning_state": "enabled",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405778 + | "Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason for failure}. "; |
| 92405846 + | "Failed to modify the object store volume. Reason: {Reason for failure}."; |
| 92405811 + | "Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |
| 92405858 + | "Failed to "modify" the "bucket" because the operation is only supported on data SVMs."; |
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92405863 + | "An error occurs when creating an access policy. The reason for failure is detailed in the error message."; |
| 92405864 + | "An error occurs when deleting an access policy. The reason for failure is detailed in the error message."; |
| 92405891 + | The resources specified in the access policy are not valid. Valid ways to specify a resource are *, <bucket-name>, <bucket-name>/…/…. Valid characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &, @,?, (, ), single quote, *, !, - and $. |
| 92405894 + | "Statements, principals and resources list can have a maximum of 10 entries."; |

| Error code | Message |
| --- | --- |
| 92405897 + | The principals specified in the access policy are not in the correct format. User name must be in between 1 and 64 characters. Valid characters for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - . |
| 92405898 + | "The SID specified in the access policy is not valid."; |
| 92406014 + | "Failed to modify event selector for bucket "{bucket name}". If the value of either access or permission is set to none, they both must be set to none."; |
| 92733458 + | "[Job job number] Job failed: Failed to modify bucket "s3bucket1" for SVM "vs1". Reason: {Reason for failure}. "; |
| 8454236 + | "Could not assign qtree "qtree1" to QoS policy group "group1". Invalid QoS policy group specified "group1". The specified QoS policy group has a min-throughput value set, and the workload being assigned resides on a platform that does not support min-throughput or the cluster is in a mixed version state and the effective cluster version of ONTAP does not support min-throughput on this platform."; |
| 8454323 + | "Policy group with UUID "23bwegew-8eqg-121r-bjad-0050e628wq732" does not exist." |
| 92406230 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be greater than the maximum lock retention period set in the object store server for SVM "{SVM}". Check the maximum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406236 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be less than the minimum lock retention period set in the object store server for SVM "{SVM}". Check the minimum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406217 + | "The specified "allowed_headers" is not valid because it contains more than one wild card ("*") character."; |
| 92406224 + | "A Cross-Origin Resource Sharing (CORS) rule must have an origin and HTTP method specified."; |
| 92406222 + | "Cannot specify Cross-Origin Resource Sharing (CORS) configuration for object store bucket "{bucket}" on SVM "{SVM}". Specifying such configuration is supported on object store volumes created in ONTAP 9.8 or later releases only."; |
| 92406211 + | "The specified method "DONE" is not valid. Valid methods are GET, PUT, DELETE, HEAD, and POST."; |

| Error code | Message |
| --- | --- |
| 92405863 + | "Failed to create CORS rules for bucket "bb1". Reason: "Field "index" cannot be specified for this operation.". Resolve all the issues and retry the operation."; |
| 92406228 + | "Cannot exceed the maximum limit of 100 Cross-Origin Resource Sharing (CORS) rules per S3 bucket "{bucket}" in SVM "{SVM}".";; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
| --- | --- | --- |
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage protocols S3 services

### Manage protocols S3 services

#### Overview

An S3 server is an object store server that is compatible with the Amazon S3 protocol. In the initial version, only a subset of the protocol features necessary to support Fabric Pool capacity tier usecases are implemented. S3 server allows you to store objects in ONTAP using Amazon S3 protocol. This feature can be used as a target object store server for ONTAP FabricPools.

#### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

### Retrieving all of the S3 configurations

```
# The API:
/api/protocols/s3/services

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/services?fields=*&return_records=true&return_timeout=
15" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "cf90b8f2-8071-11e9-8190-0050568eae21",
      "name": "vs2"
    },
    "name": "vs1",
```

```
    "comment": "S3 server",
    "enabled": false,
    "max_lock_retention_period": "none",
    "min_lock_retention_period": "none"
  },
  {
    "svm": {
      "uuid": "d7f1219c-7f8e-11e9-9124-0050568eae21",
      "name": "vs1"
    },
    "name": "Server-1",
    "comment": "S3 server",
    "enabled": true,
    "max_lock_retention_period": "none",
    "min_lock_retention_period": "none",
    "buckets": [
      {
        "uuid": "e08665af-8114-11e9-8190-0050568eae21",
        "name": "bucket-1",
        "volume": {
          "name": "fg_oss_1559026220",
          "uuid": "de146bff-8114-11e9-8190-0050568eae21"
        },
        "size": 107374182400,
        "logical_used_size": 157286400,
        "encryption": {
          "enabled": false
        },
        "comment": "s3 bucket"
      },
      {
        "uuid": "fb1912ef-8114-11e9-8190-0050568eae21",
        "name": "bucket-2",
        "volume": {
          "name": "fg_oss_1559026269",
          "uuid": "f9b1cdd0-8114-11e9-8190-0050568eae21"
        },
        "size": 107374182400,
        "logical_used_size": 78643200,
        "encryption": {
          "enabled": false
        },
        "comment": "s3 bucket"
      }
    ],
    "users": [
```

```
      {
        "name": "user-1",
        "comment": "S3 user",
        "access_key": "<AWS-ACCESS-KEY-ID>"
      },
      {
        "name": "user-2",
        "comment": "",
        "access_key": "<AWS-ACCESS-KEY-ID>"
      }
    ]
  }
],
"num_records": 2
}
```

**Retrieving all S3 configurations for a particular SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/24c2567a-f269-
11e8-8852-0050568e5298?fields=*" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "d7f1219c-7f8e-11e9-9124-0050568eae21",
  "name": "vs1"
},
"name": "Server-1",
"comment": "S3 server",
"enabled": true,
"max_lock_retention_period": "none",
"min_lock_retention_period": "none",
"buckets": [
  {
    "uuid": "e08665af-8114-11e9-8190-0050568eae21",
    "name": "bucket-1",
    "volume": {
      "name": "fg_oss_1559026220",
      "uuid": "de146bff-8114-11e9-8190-0050568eae21"
    },
    "size": 107374182400,
```

```
    "logical_used_size": 157286400,
    "encryption": {
      "enabled": false
    },
    "comment": "s3 bucket",
    "policy": {
      "statements": [
        {
          "effect": "deny",
          "actions": [
            "*Object"
          ],
          "principals": [
            "mike"
          ],
          "resources": [
            "bucket-1/policy-docs/*",
            "bucket-1/confidential-*"
          ],
          "sid": "DenyAccessToGetPutDeleteObjectForMike"
        },
        {
          "effect": "allow",
          "actions": [
            "GetObject"
          ],
          "principals": [
            "*"
          ],
          "resources": [
            "bucket-1/readme"
          ],
          "sid": "AccessToGetObjectForAnonymousUser"
        }
      ]
    },
    "cors": {
      "rules": [
        {
          "id": "string",
          "allowed_origins": [
            "http://www.example.com"
          ],
          "allowed_methods": [
            "PUT",
            "DELETE"
```

```
          ],
          "allowed_headers": [
            "x-amz-request-id"
          ],
          "expose_headers": [
            "http://www.example.com"
          ],
          "max_age_seconds": 1024
        }
      ]
    }
  },
  {
    "uuid": "fb1912ef-8114-11e9-8190-0050568eae21",
    "name": "bucket-2",
    "volume": {
      "name": "fg_oss_1559026269",
      "uuid": "f9b1cdd0-8114-11e9-8190-0050568eae21"
    },
    "size": 107374182400,
    "logical_used_size": 1075838976,
    "encryption": {
      "enabled": false
    },
    "comment": "s3 bucket"
  }
],
"users": [
  {
    "name": "user-1",
    "comment": "s3 user",
    "access_key": "<AWS-ACCESS-KEY-ID>"
  },
  {
    "name": "user-2",
    "comment": "",
    "access_key": "<AWS-ACCESS-KEY-ID>"
  }
]
}
```


**Creating an S3 server, users, and buckets configurations with required fields specified**

```
# The API:
/api/protocols/s3/services
```

```
# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"buckets\": [ { \"name\": \"bucket-1\" }, { \"name\": \"bucket-2\" } ],
\"enabled\": true, \"name\": \"Server-1\", \"svm\": { \"uuid\":
\"d49ef663-7f8e-11e9-9b2c-0050568e4594\" }, \"users\": [ { \"name\":
\"user-1\" }, { \"name\": \"user-2\" } ]}"

# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 08:44:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/
Content-Length: 623
Content-Type: application/hal+json
{
"num_records": 1,
"records": [
  {
    "users": [
      {
        "name": "user-1",
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "secret_key": "<AWS-SECRET-ACCESS-KEY>"
      },
      {
        "name": "user-2",
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "secret_key": "<AWS-SECRET-ACCESS-KEY>"
      }
    ],
    "job": {
      "uuid": "f51675dd-820a-11e9-a762-0050568e4594",
      "_links": {
        "self": {
          "href": "/api/cluster/jobs/f51675dd-820a-11e9-a762-0050568e4594"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/"
      }
```

```
    }
  }
]
}
```

### Creating an S3 server, users, and buckets configurations

```
# The API:
/api/protocols/s3/services

# The call:

curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"buckets\": [ { \"aggregates\": [ { \"name\": \"aggr1\", \"uuid\":
\"1cd8a442-86d1-11e0-ae1c-123478563412\" } ],
\"constituents_per_aggregate\": 4, \"name\": \"bucket-1\", \"size\":
\"107374182400\",  \"policy\": { \"statements\": [ { \"actions\": [ \"*\"
], \"conditions\": [ { \"operator\": \"ip_address\", \"source_ips\": [
\"1.1.1.1/23\", \"1.2.2.2/20\" ] } ], \"effect\": \"allow\",
\"resources\": [ \"bucket-1\", \"bucket-1*\" ], \"sid\":
\"fullAccessForAllPrincipalsToBucket\"} ] } }, { \"aggregates\": [ {
\"name\": \"aggr1\", \"uuid\": \"1cd8a442-86d1-11e0-ae1c-123478563412\" },
{ \"name\": \"aggr2\", \"uuid\": \"982fc4d0-d1a2-4da4-9c47-5b433f24757d\"}
], \"constituents_per_aggregate\": 4, \"name\": \"bucket-2\" } ],
\"enabled\": true, \"name\": \"Server-1\", \"max_lock_retention_period\":
\"P41Y\", \"max_lock_retention_period\": \"P1Y\", \"svm\": { \"name\":
\"vs1\", \"uuid\": \"d49ef663-7f8e-11e9-9b2c-0050568e4594\" }, \"users\":
[ { \"name\": \"user-1\" }, { \"name\": \"user-2\" } ]}"



# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 08:44:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/
Content-Length: 623
Content-Type: application/hal+json
{
"num_records": 1,
"records": [
  {
```

```
    "users": [
      {
        "name": "user-1",
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "secret_key": "<AWS-SECRET-ACCESS-KEY>"
      },
      {
        "name": "user-2",
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "secret_key": "<AWS-SECRET-ACCESS-KEY>"
      }
    ],
    "job": {
      "uuid": "f51675dd-820a-11e9-a762-0050568e4594",
      "_links": {
        "self": {
          "href": "/api/cluster/jobs/f51675dd-820a-11e9-a762-0050568e4594"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/"
      }
    }
  }
]
}
```


**Creating an S3 server configuration (HTTPS)**

```
# The API:
/api/protocols/s3/services

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"comment\":
\"http S3 server\", \"enabled\": true, \"name\": \"Server-1\", \"svm\": {
\"name\": \"vs1\", \"uuid\": \"db2ec036-8375-11e9-99e1-0050568e3ed9\" },
\"is_https_enabled\": \"false\" }"
```

**Creating an S3 server configuration (HTTP)**

```
# The API:
/api/protocols/s3/services

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"comment\":
\"https S3 server\", \"enabled\": true, \"name\": \"Server-1\", \"svm\": {
\"name\": \"vs1\", \"uuid\": \"db2ec036-8375-11e9-99e1-0050568e3ed9\" },
\"certificate\": { \"uuid\": \"db2ec036-8375-11e9-99e1-0050568e3ed9\" }}"
```

**Disable s3 server for the specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"enabled\": false }"
```

**Deleting the S3 server for a specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/a425f10b-ad3b-
11e9-b559-0050568e8222?delete_all=false" -H  "accept: application/json"
HTTP/1.1 200 OK
Date: Wed, 14 Aug 2019 07:04:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 132
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "job": {
      "uuid": "bf74ba50-be61-11e9-bea8-0050568e8222"
    }
  }
]
}
```

**Deleting all of the S3 server configuration for a specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298?delete_all=true" -H "accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Sat, 01 Jun 2019 15:46:39 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "71eaaf02-8484-11e9-91f7-0050568ebc5f",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/71eaaf02-8484-11e9-91f7-0050568ebc5f"
    }
  }
}
}
```

### Retrieve the S3 server configuration for all SVMs

GET `/protocols/s3/services`

**Introduced In:** 9.7

Retrieves the S3 server configuration for all SVMs. Note that in order to retrieve S3 bucket policy conditions, 'fields' option should be set to '**'.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver object-store-server show`

**Learn more**

- DOC /protocols/s3/services

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br>• Introduced in: 9.8 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br>• Introduced in: 9.8 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br>• Introduced in: 9.8 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br>• Introduced in: 9.8 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br>• Introduced in: 9.8 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.8 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.8 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.8 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name<br>• maxLength: 253<br>• minLength: 3 |
| is_http_enabled | boolean | query | False | Filter by is_http_enabled<br>• Introduced in: 9.8 |
| users.comment | string | query | False | Filter by users.comment<br>• maxLength: 256<br>• minLength: 0 |
| users.key_time_to_live | string | query | False | Filter by users.key_time_to_live<br>• Introduced in: 9.14 |
| users.svm.name | string | query | False | Filter by users.svm.name |
| users.svm.uuid | string | query | False | Filter by users.svm.uuid |
| users.name | string | query | False | Filter by users.name<br>• maxLength: 64<br>• minLength: 1 |
| users.key_expiry_time | string | query | False | Filter by users.key_expiry_time<br>• Introduced in: 9.14 |
| users.access_key | string | query | False | Filter by users.access_key |
| enabled | boolean | query | False | Filter by enabled |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.8 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.8 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.8 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.8 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.8 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.8 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.8 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br>• Introduced in: 9.8 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br>• Introduced in: 9.8 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.8 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.8 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.8 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.8 |
| default_unix_user | string | query | False | Filter by default_unix_user<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| port | integer | query | False | Filter by port<br>• Introduced in: 9.8<br>• Max value: 65535<br>• Min value: 1 |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| default_win_user | string | query | False | Filter by default_win_user<br>• Introduced in: 9.12 |
| max_key_time_to_live | string | query | False | Filter by max_key_time_to_live<br>• Introduced in: 9.15 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| buckets.is_nas_path_mutable | boolean | query | False | Filter by buckets.is_nas_path_mutable<br>• Introduced in: 9.17 |
| buckets.nas_path | string | query | False | Filter by buckets.nas_path<br>• Introduced in: 9.12 |
| buckets.encryption.enabled | boolean | query | False | Filter by buckets.encryption.enabled |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.allowed | boolean | query | False | Filter by buckets.allowed<br><br>• Introduced in: 9.12 |
| buckets.protection_status.destination.is_ontap | boolean | query | False | Filter by buckets.protection_status.destination.is_ontap<br><br>• Introduced in: 9.10 |
| buckets.protection_status.destination.is_cloud | boolean | query | False | Filter by buckets.protection_status.destination.is_cloud<br><br>• Introduced in: 9.10 |
| buckets.protection_status.destination.is_external_cloud | boolean | query | False | Filter by buckets.protection_status.destination.is_external_cloud<br><br>• Introduced in: 9.12 |
| buckets.protection_status.is_protected | boolean | query | False | Filter by buckets.protection_status.is_protected<br><br>• Introduced in: 9.10 |
| buckets.svm.name | string | query | False | Filter by buckets.svm.name |
| buckets.svm.uuid | string | query | False | Filter by buckets.svm.uuid |
| buckets.retention.default_period | string | query | False | Filter by buckets.retention.default_period<br><br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.retention.mode | string | query | False | Filter by buckets.retention.mode<br>• Introduced in: 9.14 |
| buckets.role | string | query | False | Filter by buckets.role<br>• Introduced in: 9.10 |
| buckets.size | integer | query | False | Filter by buckets.size<br>• Max value: 67553994410557440<br>• Min value: 107374182400 |
| buckets.logical_used_size | integer | query | False | Filter by buckets.logical_used_size |
| buckets.uuid | string | query | False | Filter by buckets.uuid |
| buckets.volume.name | string | query | False | Filter by buckets.volume.name |
| buckets.volume.uuid | string | query | False | Filter by buckets.volume.uuid |
| buckets.comment | string | query | False | Filter by buckets.comment<br>• maxLength: 256<br>• minLength: 0 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| buckets.qos_policy.min_throughput_mbps | integer | query | False | Filter by buckets.qos_policy.min_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| buckets.qos_policy.max_throughput_iops | integer | query | False | Filter by buckets.qos_policy.max_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| buckets.qos_policy.max_throughput_mbps | integer | query | False | Filter by buckets.qos_policy.max_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| buckets.qos_policy.name | string | query | False | Filter by buckets.qos_policy.name<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.qos_policy.min_throughput_iops | integer | query | False | Filter by buckets.qos_policy.min_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| buckets.qos_policy.uuid | string | query | False | Filter by buckets.qos_policy.uuid<br>• Introduced in: 9.8 |
| buckets.qos_policy.max_throughput | string | query | False | Filter by buckets.qos_policy.max_throughput<br>• Introduced in: 9.17 |
| buckets.qos_policy.min_throughput | string | query | False | Filter by buckets.qos_policy.min_throughput<br>• Introduced in: 9.17 |
| buckets.versioning_state | string | query | False | Filter by buckets.versioning_state<br>• Introduced in: 9.11 |
| buckets.is_consistent_etag | boolean | query | False | Filter by buckets.is_consistent_etag<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.audit_event_selector.permission | string | query | False | Filter by buckets.audit_event_selector.permission<br>• Introduced in: 9.10 |
| buckets.audit_event_selector.access | string | query | False | Filter by buckets.audit_event_selector.access<br>• Introduced in: 9.10 |
| buckets.type | string | query | False | Filter by buckets.type<br>• Introduced in: 9.12 |
| buckets.lifecycle_management.rules.bucket_name | string | query | False | Filter by buckets.lifecycle_management.rules.bucket_name<br>• Introduced in: 9.14<br>• maxLength: 63<br>• minLength: 3 |
| buckets.lifecycle_management.rules.non_current_version_expiration.non_current_days | integer | query | False | Filter by buckets.lifecycle_management.rules.non_current_version_expiration.non_current_days<br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.non_current_version_expiration.new_non_current_versions | integer | query | False | Filter by buckets.lifecycle_management.rules.non_current_version_expiration.new_non_current_versions<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.lifecycle_management.rules.uuid | string | query | False | Filter by buckets.lifecycle_management.rules.uuid<br>• Introduced in: 9.14 |
| buckets.lifecycle_management.rules.expiration.expired_object_delete_marker | boolean | query | False | Filter by buckets.lifecycle_management.rules.expiration.expired_object_delete_marker<br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.expiration.object_expiry_date | string | query | False | Filter by buckets.lifecycle_management.rules.expiration.object_expiry_date<br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.expiration.object_age_days | integer | query | False | Filter by buckets.lifecycle_management.rules.expiration.object_age_days<br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.name | string | query | False | Filter by buckets.lifecycle_management.rules.name<br>• Introduced in: 9.13<br>• maxLength: 256<br>• minLength: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.lifecycle_management.rules.svm.name | string | query | False | Filter by buckets.lifecycle_management.rules.svm.name<br><br>• Introduced in: 9.14 |
| buckets.lifecycle_management.rules.svm.uuid | string | query | False | Filter by buckets.lifecycle_management.rules.svm.uuid<br><br>• Introduced in: 9.14 |
| buckets.lifecycle_management.rules.object_filter.prefix | string | query | False | Filter by buckets.lifecycle_management.rules.object_filter.prefix<br><br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.object_filter.size_greater_than | integer | query | False | Filter by buckets.lifecycle_management.rules.object_filter.size_greater_than<br><br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.object_filter.size_less_than | integer | query | False | Filter by buckets.lifecycle_management.rules.object_filter.size_less_than<br><br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.object_filter.tags | string | query | False | Filter by buckets.lifecycle_management.rules.object_filter.tags<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.lifecycle_management.rules.enabled | boolean | query | False | Filter by buckets.lifecycle_management.rules.enabled<br>• Introduced in: 9.13 |
| buckets.lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days | integer | query | False | Filter by buckets.lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days<br>• Introduced in: 9.13 |
| buckets.policy.statements.principals | string | query | False | Filter by buckets.policy.statements.principals<br>• Introduced in: 9.8 |
| buckets.policy.statements.sid | string | query | False | Filter by buckets.policy.statements.sid<br>• Introduced in: 9.8<br>• maxLength: 256<br>• minLength: 0 |
| buckets.policy.statements.resources | string | query | False | Filter by buckets.policy.statements.resources<br>• Introduced in: 9.8 |
| buckets.policy.statements.effect | string | query | False | Filter by buckets.policy.statements.effect<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.policy.statements.actions | string | query | False | Filter by buckets.policy.statements.actions<br><br>• Introduced in: 9.8 |
| buckets.policy.statements.conditions.operator | string | query | False | Filter by buckets.policy.statements.conditions.operator<br><br>• Introduced in: 9.8 |
| buckets.policy.statements.conditions.delimiters | string | query | False | Filter by buckets.policy.statements.conditions.delimiters<br><br>• Introduced in: 9.8 |
| buckets.policy.statements.conditions.prefixes | string | query | False | Filter by buckets.policy.statements.conditions.prefixes<br><br>• Introduced in: 9.8 |
| buckets.policy.statements.conditions.usernames | string | query | False | Filter by buckets.policy.statements.conditions.usernames<br><br>• Introduced in: 9.8 |
| buckets.policy.statements.conditions.source_ips | string | query | False | Filter by buckets.policy.statements.conditions.source_ips<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| buckets.policy.statements.conditions.max_keys | integer | query | False | Filter by buckets.policy.statements.conditions.max_keys<br>• Introduced in: 9.8 |
| buckets.name | string | query | False | Filter by buckets.name<br>• maxLength: 63<br>• minLength: 3 |
| buckets.snapshot_restore.objects_remaining | integer | query | False | Filter by buckets.snapshot_restore.objects_remaining<br>• Introduced in: 9.18 |
| buckets.snapshot_restore.state | string | query | False | Filter by buckets.snapshot_restore.state<br>• Introduced in: 9.18 |
| buckets.snapshot_restore.snapshot | string | query | False | Filter by buckets.snapshot_restore.snapshot<br>• Introduced in: 9.18 |
| buckets.snapshot_restore.progress | integer | query | False | Filter by buckets.snapshot_restore.progress<br>• Introduced in: 9.18 |
| buckets.snapshot_policy.uuid | string | query | False | Filter by buckets.snapshot_policy.uuid<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| buckets.snapshot_policy.name | string | query | False | Filter by buckets.snapshot_policy.name<br>• Introduced in: 9.16 |
| buckets.cors.rules.allowed_headers | string | query | False | Filter by buckets.cors.rules.allowed_headers<br>• Introduced in: 9.16 |
| buckets.cors.rules.id | string | query | False | Filter by buckets.cors.rules.id<br>• Introduced in: 9.16<br>• maxLength: 256<br>• minLength: 0 |
| buckets.cors.rules.max_age_seconds | integer | query | False | Filter by buckets.cors.rules.max_age_seconds<br>• Introduced in: 9.16 |
| buckets.cors.rules.allowed_origins | string | query | False | Filter by buckets.cors.rules.allowed_origins<br>• Introduced in: 9.16 |
| buckets.cors.rules.expose_headers | string | query | False | Filter by buckets.cors.rules.expose_headers<br>• Introduced in: 9.16 |
| buckets.cors.rules.allowed_methods | string | query | False | Filter by buckets.cors.rules.allowed_methods<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| min_lock_retention_period | string | query | False | Filter by min_lock_retention_period<br>• Introduced in: 9.16 |
| max_lock_retention_period | string | query | False | Filter by max_lock_retention_period<br>• Introduced in: 9.16 |
| secure_port | integer | query | False | Filter by secure_port<br>• Introduced in: 9.8<br>• Max value: 65535<br>• Min value: 1 |
| certificate.name | string | query | False | Filter by certificate.name<br>• Introduced in: 9.8 |
| certificate.uuid | string | query | False | Filter by certificate.uuid<br>• Introduced in: 9.8 |
| is_https_enabled | boolean | query | False | Filter by is_https_enabled<br>• Introduced in: 9.8 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "buckets": [
        {
          "audit_event_selector": {
            "access": "string",
            "permission": "string"
          },
          "comment": "S3 bucket.",
          "cors": {
            "rules": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "allowed_headers": [
                  "x-amz-request-id"
                ],
                "allowed_methods": [
                  "PUT",
                  "DELETE"
                ],
                "allowed_origins": [
                  "http://www.example.com"
                ],
                "expose_headers": [
                  "x-amz-date"
```

```
                ],
                "id": "string",
                "max_age_seconds": 1024
              }
            ]
          },
          "lifecycle_management": {
            "rules": [
              {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "abort_incomplete_multipart_upload": {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  }
                },
                "bucket_name": "bucket1",
                "expiration": {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "object_age_days": 100,
                  "object_expiry_date": "2039-09-23 00:00:00 +0000"
                },
                "name": "string",
                "non_current_version_expiration": {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  }
                },
                "object_filter": {
                  "_links": {
                    "self": {
                      "href": "/api/resourcelink"
                    }
                  },
                  "prefix": "/logs",
```

```
                "size_greater_than": 10240,
                "size_less_than": 10485760,
                "tags": [
                  "project1=projA",
                  "project2=projB"
                ]
              },
              "svm": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "svm1",
                "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
              },
              "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
            }
          ]
        },
        "logical_used_size": 0,
        "name": "bucket1",
        "nas_path": "/",
        "policy": {
          "statements": [
            {
              "actions": [
                "GetObject",
                "PutObject",
                "DeleteObject",
                "ListBucket"
              ],
              "conditions": [
                {
                  "delimiters": [
                    "/"
                  ],
                  "max_keys": [
                    1000
                  ],
                  "operator": "ip_address",
                  "prefixes": [
                    "pref"
                  ],
                  "source_ips": [
                    "1.1.1.1",
```

```
                "1.2.2.0/24"
              ],
              "usernames": [
                "user1"
              ]
            }
          ],
          "effect": "allow",
          "principals": [
            "user1",
            "group/grp1",
            "nasgroup/group1"
          ],
          "resources": [
            "bucket1",
            "bucket1/*"
          ],
          "sid": "Full_Access_To_User1!"
        }
      ]
    },
    "qos_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "max_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
```

```
            ],
            "min_throughput_iops": 2000,
            "min_throughput_mbps": 500,
            "name": "performance",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "retention": {
            "default_period": "P10Y",
            "mode": "governance"
          },
          "role": "string",
          "size": 214748364800,
          "snapshot_policy": {
            "name": "default-1weekly",
            "uuid": "3675af31-431c-12fa-114a-20675afebc12"
          },
          "snapshot_restore": {
            "snapshot": "string",
            "state": "string"
          },
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "type": "s3",
          "use_mirrored_aggregates": true,
          "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
          "versioning_state": "enabled",
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "certificate": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "comment": "S3 server",
      "default_unix_user": "string",
      "default_win_user": "string",
      "max_key_time_to_live": "PT6H3M",
      "max_lock_retention_period": "P10Y",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "min_lock_retention_period": "P10Y",
      "name": "Server-1",
      "port": 80,
      "secure_port": 443,
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
```

```
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        {
          "access_key": "<AWS-ACCESS-KEY-ID>",
          "comment": "S3 user",
          "key_expiry_time": "2024-01-01 00:00:00 +0000",
          "key_time_to_live": "PT6H3M",
          "name": "user-1",
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          }
        }
      ]
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
|---|---|---|
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |

| Name | Type | Description |
| --- | --- | --- |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| svm | svm | SVM, applies only to SVM-scoped objects. |

s3_service

Specifies the S3 server configuration.

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| buckets | array[s3_bucket] | This field cannot be specified in a PATCH method. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |

| Name | Type | Description |
| --- | --- | --- |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |

| Name | Type | Description |
|---|---|---|
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[s3_user] | This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create S3 server, users, and buckets configurations**

POST `/protocols/s3/services`

**Introduced In:** 9.7

Creates an S3 server, users, and buckets configurations.

#### Important notes

- Each SVM can have one S3 server configuration.
- By default, HTTPS is enabled on the S3 server, so a valid certificate must be provided when creating a S3 Server for the SVM.
- One or more buckets and users can also be created using this end-point.
- If creating a user configuration fails, buckets are not created either and already created users are not saved.
- If creating a bucket configuration fails, all buckets already created are saved with no new buckets created.

#### Required properties

- `svm.uuid` - Existing SVM in which to create an S3 server configuration.

#### Recommended optional properties

- `enabled` - Specifies the state of the server created.
- `comment` - Any information related to the server created.

#### Default property values

- `comment` - ""
- `enabled` - *true*
- `is_https_enabled` - *true*

#### Related ONTAP commands

- `vserver object-store-server create`
- `vserver object-store-server bucket create`
- `vserver object-store-server bucket policy statement create`
- `vserver object-store-server bucket policy-statement-condition create`
- `vserver object-store-server bucket cors-rule create`
- `vserver object-store-server user create`

#### Learn more

- DOC /protocols/s3/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| buckets | array[s3_bucket] | This field cannot be specified in a PATCH method. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |

| Name | Type | Description |
| --- | --- | --- |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[s3_user] | This field cannot be specified in a PATCH method. |

**Example request**

```
{
  "buckets": [
    {
      "aggregates": [
        {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "audit_event_selector": {
        "access": "string",
        "permission": "string"
      },
      "comment": "S3 bucket.",
      "constituents_per_aggregate": 4,
      "cors": {
        "rules": [
          {
            "allowed_headers": [
              "x-amz-request-id"
            ],
            "allowed_methods": [
              "PUT",
              "DELETE"
            ],
            "allowed_origins": [
              "http://www.example.com"
            ],
            "expose_headers": [
              "x-amz-date"
            ],
            "id": "string",
            "max_age_seconds": 1024
          }
        ]
      },
      "lifecycle_management": {
        "rules": [
          {
            "bucket_name": "bucket1",
            "expiration": {
              "object_age_days": 100,
              "object_expiry_date": "2039-09-23 00:00:00 +0000"
            },
```

```
            "name": "string",
            "object_filter": {
              "prefix": "/logs",
              "size_greater_than": 10240,
              "size_less_than": 10485760,
              "tags": [
                "project1=projA",
                "project2=projB"
              ]
            },
            "svm": {
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
          }
        ]
      },
      "logical_used_size": 0,
      "name": "bucket1",
      "nas_path": "/",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  1000
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
```

```
                "usernames": [
                  "user1"
                ]
              }
            ],
            "effect": "allow",
            "principals": [
              "user1",
              "group/grp1",
              "nasgroup/group1"
            ],
            "resources": [
              "bucket1",
              "bucket1/*"
            ],
            "sid": "Full_Access_To_User1!"
          }
        ]
      },
      "qos_policy": {
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "retention": {
```

```
        "default_period": "P10Y",
        "mode": "governance"
      },
      "role": "string",
      "size": 214748364800,
      "snapshot_policy": {
        "name": "default-1weekly",
        "uuid": "3675af31-431c-12fa-114a-20675afebc12"
      },
      "snapshot_restore": {
        "snapshot": "string",
        "state": "string"
      },
      "storage_service_level": "value",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "s3",
      "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
      "versioning_state": "enabled",
      "volume": {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "comment": "S3 server",
  "default_unix_user": "string",
  "default_win_user": "string",
  "max_key_time_to_live": "PT6H3M",
  "max_lock_retention_period": "P10Y",
  "min_lock_retention_period": "P10Y",
  "name": "Server-1",
  "port": 80,
  "secure_port": 443,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
```

```
      "access_key": "<AWS-ACCESS-KEY-ID>",
      "comment": "S3 user",
      "key_expiry_time": "2024-01-01 00:00:00 +0000",
      "key_time_to_live": "PT6H3M",
      "name": "user-1",
      "secret_key": "<AWS-SECRET-ACCESS-KEY>",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of Records |
| records | array[records] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "job": {
        "uuid": "string"
      },
      "users": [
        {
          "access_key": "<AWS-ACCESS-KEY-ID>",
          "key_expiry_time": "2024-01-01 00:00:00 +0000",
          "name": "user-1",
          "secret_key": "<AWS-SECRET-ACCESS-KEY>"
        }
      ]
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115127 + | The cluster lacks a valid S3 license. |
| 2621706 + | The specified "{svm.uuid}" and "{svm.name}" refer to different SVMs. |
| 92405789 + | The specified object server name contains invalid characters or not a fully qualified domain name. Valid characters for an object store server name are 0-9, A-Z, a-z, ".", and "-". |
| 92405790 + | Object store server names must have between 3 and 253 characters. |

| Error Code | Description |
| --- | --- |
| 92405839 + | Creating an object store server requires an effective cluster version of data ONTAP 9.7.0 or later. Upgrade all the nodes to 9.7.0 or later and try the operation again. |
| 92405853 + | Failed to create the object store server because Cloud Volumes ONTAP does not support object store servers. |
| 92405863 + | An error occurs when creating an S3 user or bucket. The reason for failure is detailed in the error message. Follow the error codes specified for the user or bucket endpoints to see details for the failure. |
| 92405863 + | Failed to create bucket "{bucket name}". Reason: "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason: Bucket name "{bucket name}" contains invalid characters or invalid character combinations. Valid characters for a bucket name are 0-9, a-z, ".", and "-". Invalid character combinations are ".-", "-.", and "..". ". Resolve all the issues and retry the operation. |
| 92405863 + | Failed to create bucket "{bucket name}". Reason: "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason: Invalid QoS policy group specified "{qos policy}". The specified QoS policy group has a min-throughput value set, and the workload being assigned resides on a platform that does not support min-throughput or the cluster is in a mixed version state and the effective cluster version of ONTAP does not support min-throughput on this platform. Resolve all the issues and retry the operation. |
| 92405863 + | Failed to create bucket "{bucket name}". Reason: "User(s) "{user name(s)}" specified in the principal list do not exist for SVM "{svm.name}". Use the "object-store-server user create" command to create a user.". Resolve all the issues and retry the operation. |
| 92405863 + | Failed to create user "{user name}". Reason: "SVM "Cluster" is not a data SVM. Specify a data SVM.". Resolve all the issues and retry the operation. |
| 92405884 + | An object store server can only be created on a data SVM. An object store server can also be created on a system SVM on a mixed platform cluster. |
| 92405903 + | Failed to configure HTTPS on an object store server for SVM "{svm.name}". Reason: {Reason of failure}. |
| 92405900 + | Certificate not found for SVM "{svm.name}". |
| 92405917 + | The specified certificate name and UUID do not refer to the same certificate. |

| Error Code | Description |
| --- | --- |
| 92406020 + | Only certificates of type "server" are supported. |
| 92406044 + | Failed to set default UNIX user for SVM "{svm.name}". Reason: UNIX user can only be created on a Data SVM. |
| 92406071 + | S3 protocol is not present in the allowed protocol list for SVM "{svm.name}". |
| 92406196 + | The specified value for the "key_time_to_live" field cannot be greater than the maximum limit specified for the "max_key_time_to_live" field in the object store server. |
| 92406197 + | Object store user "user-2" must have a non-zero value for the "key_time_to_live" field because the maximum limit specified for the "max_key_time_to_live" field in the object store server is not zero. |
| 92406230 + | The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be greater than the maximum lock retention period set in the object store server for SVM "{SVM}". Check the maximum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again. |
| 92406231 + | One or more object store buckets exist with a default retention period greater than the "max_lock_retention_period" specified. Check the default retention period set for each bucket in the specified SVM and try the operation again. |
| 92406236 + | The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be less than the minimum lock retention period set in the object store server for SVM "{SVM}". Check the minimum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again. |
| 92406237 + | One or more object store buckets exist with a default retention period less than the "min_lock_retention_period" specified. Check the default retention period set for each bucket in the specified SVM and try the operation again. |
| 92406238 + | The value for the "min_lock_retention_period" parameter cannot be greater than the "max_lock_retention_period" parameter for the object store server for SVM "vs1". |
| 92406217 + | The specified "-allowed-headers" in not valid because it contains more than one wild card ("__") character.; |
| 92406224 + | A Cross-Origin Resource Sharing (CORS) rule must have an origin and HTTP method specified.; |

| Error Code | Description |
| --- | --- |
| 92406211 + | The specified method "DONE" is not valid. Valid methods are GET, PUT, DELETE, HEAD, and POST.; |
| 92405863 + | Failed to create CORS rules for bucket "bb1". Reason: "Field "index" cannot be specified for this operation.". Resolve all the issues and retry the operation.; |
| 92406228 + | Cannot exceed the maximum limit of 100 Cross-Origin Resource Sharing (CORS) rules per S3 bucket "{bucket}" in SVM "{SVM}".; |
| 92405968 + | Subnet "{condition.source_ips}" is not a valid IP subnet because it contains non-zero values in the host component of the address. Subnet address "{valid source ips subnet}" identifies a valid subnet for the given mask length. Example: "10.0.1.0/24" is a valid subnet while as"10.0.1.1/24" is invalid."; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |

| Name | Type | Description |
| --- | --- | --- |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
|---|---|---|
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume. |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| constituents_per_aggregate | integer | Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |

| Name | Type | Description |
| --- | --- | --- |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| storage_service_level | string | Specifies the storage service level of the FlexGroup volume on which the bucket should be created. Valid values are "value", "performance" or "extreme". |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup. Only mirrored aggregates are used if this parameter is set to "true" and only unmirrored aggregates are used if this parameter is set to "false". The default value is "true" for a MetroCluster configuration and is "false" for a non-MetroCluster configuration. |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

s3_service

Specifies the S3 server configuration.

| Name | Type | Description |
| --- | --- | --- |
| buckets | array[s3_bucket] | This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |

| Name | Type | Description |
| --- | --- | --- |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |

| Name | Type | Description |
| --- | --- | --- |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[s3_user] | This field cannot be specified in a PATCH method. |

collection_links

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

s3_service_user_post_response

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| key_expiry_time | string | Specifies the date and time after which the keys expire and are no longer valid. |
| name | string | The name of the user. |
| secret_key | string | Specifies the secret key for the user. |

warning

Specifies a warning message sent from the S3 server during a POST or PATCH operation.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Warning code of the warning encountered. |

| Name | Type | Description |
|---|---|---|
| message | string | Details of the warning sent from the S3 server. |

records

| Name | Type | Description |
|---|---|---|
| job | job_link | |
| users | array[s3_service_user_post_response] | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete the S3 server configuration for an SVM

DELETE `/protocols/s3/services/{svm.uuid}`

**Introduced In:** 9.7

Deletes the S3 server configuration of an SVM. If the 'delete_all' parameter is set to false, only the S3 server is deleted. Otherwise S3 users and buckets present on the SVM are also deleted. Note that only empty buckets can be deleted. This endpoint returns the S3 server delete job-uuid in response. To monitor the job status follow /api/cluster/jobs/<job-uuid>.</job-uuid>

#### Related ONTAP commands

- `vserver object-store-server delete`

**Learn more**

- DOC /protocols/s3/services

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| delete_all | boolean | query | False | Delete S3 server and associated users and empty buckets.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405864 | An error occurs when deleting an S3 user or bucket. The reason for failure is detailed in the error message. Follow the error codes specified for the user or bucket endpoints to see details for the failure. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve the S3 server configuration for an SVM**

GET `/protocols/s3/services/{svm.uuid}`

**Introduced In:** 9.7

Retrieves the S3 Server configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server show`

#### Learn more

- DOC /protocols/s3/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| buckets | array[s3_bucket] | This field cannot be specified in a PATCH method. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |

| Name | Type | Description |
| --- | --- | --- |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[s3_user] | This field cannot be specified in a PATCH method. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "buckets": [
    {
      "audit_event_selector": {
        "access": "string",
        "permission": "string"
      },
      "comment": "S3 bucket.",
      "cors": {
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "allowed_headers": [
              "x-amz-request-id"
            ],
            "allowed_methods": [
              "PUT",
              "DELETE"
            ],
            "allowed_origins": [
              "http://www.example.com"
            ],
            "expose_headers": [
              "x-amz-date"
            ],
            "id": "string",
            "max_age_seconds": 1024
          }
        ]
      },
      "lifecycle_management": {
        "rules": [
          {
            "_links": {
              "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "abort_incomplete_multipart_upload": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        },
        "bucket_name": "bucket1",
        "expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "name": "string",
        "non_current_version_expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        },
        "object_filter": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
                }
              },
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
          }
        ]
      },
      "logical_used_size": 0,
      "name": "bucket1",
      "nas_path": "/",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  1000
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ],
            "effect": "allow",
            "principals": [
              "user1",
              "group/grp1",
```

```
            "nasgroup/group1"
          ],
          "resources": [
            "bucket1",
            "bucket1/*"
          ],
          "sid": "Full_Access_To_User1!"
        }
      ]
    },
    "qos_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "max_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "retention": {
      "default_period": "P10Y",
      "mode": "governance"
    },
    "role": "string",
```

```
      "size": 214748364800,
      "snapshot_policy": {
        "name": "default-1weekly",
        "uuid": "3675af31-431c-12fa-114a-20675afebc12"
      },
      "snapshot_restore": {
        "snapshot": "string",
        "state": "string"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "s3",
      "use_mirrored_aggregates": true,
      "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
      "versioning_state": "enabled",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "comment": "S3 server",
  "default_unix_user": "string",
  "default_win_user": "string",
  "max_key_time_to_live": "PT6H3M",
```

```
  "max_lock_retention_period": "P10Y",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "min_lock_retention_period": "P10Y",
  "name": "Server-1",
  "port": 80,
  "secure_port": 443,
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
```

```
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "access_key": "<AWS-ACCESS-KEY-ID>",
      "comment": "S3 user",
      "key_expiry_time": "2024-01-01 00:00:00 +0000",
      "key_time_to_live": "PT6H3M",
      "name": "user-1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |

| Name | Type | Description |
| --- | --- | --- |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |

| Name | Type | Description |
| --- | --- | --- |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the S3 server configuration for an SVM

PATCH `/protocols/s3/services/{svm.uuid}`

**Introduced In:** 9.7

Updates the S3 Server configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server modify`

#### Learn more

- DOC /protocols/s3/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |

| Name | Type | Description |
| --- | --- | --- |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |

| Name | Type | Description |
| --- | --- | --- |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |

### Example request

```
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "comment": "S3 server",
  "default_unix_user": "string",
  "default_win_user": "string",
  "max_key_time_to_live": "PT6H3M",
  "max_lock_retention_period": "P10Y",
  "min_lock_retention_period": "P10Y",
  "name": "Server-1",
  "port": 80,
  "secure_port": 443
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405789 + | The name "{object server name}" is not valid. A valid object server name must be a fully qualified domain name. |
| 92405790 + | Object store server name is not valid. Object store server names must have between 3 and 253 characters. |
| 92405900 + | Certificate not found for SVM "{svm.name}". |
| 92405917 + | The specified certificate name and UUID do not refer to the same certificate. |
| 92406020 + | Only certificates of type "server" are supported. + |
| | 92406153 + |

| Error Code | Description |
| --- | --- |
| Set the enabled field of the server to "down" before modifying following fields: {field name} | 92406231 + |
| One or more object store buckets exist with a default retention period greater than the "max_lock_retention_period" specified. Check the default retention period set for each bucket in the specified SVM and try the operation again. | 92406237 + |
| One or more object store buckets exist with a default retention period less than the "min_lock_retention_period" specified. Check the default retention period set for each bucket in the specified SVM and try the operation again. | 92406238 + |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Audit event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be specified in a POST method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. This field cannot be specified using a POST method. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. This field cannot be specified using a POST method. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be specified using a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
| --- | --- | --- |
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. This field cannot be specified in a POST method. |
| audit_event_selector | audit_event_selector | Audit event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be specified using a POST or PATCH method. |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be specified using a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 100GB to 60PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". |
| uuid | string | Specifies the unique identifier of the bucket. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

s3_service

Specifies the S3 server configuration.

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| comment | string | Can contain any additional information about the server being created or modified. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |

| Name | Type | Description |
| --- | --- | --- |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether the S3 server being created or modified should be up or down. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server. |
| max_key_time_to_live | string | Indicates the maximum time period that an S3 user can specify for the 'key_time_to_live' property.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If no value is specified for this property or the value specified is '0' seconds, then a user can specify any valid value. |

| Name | Type | Description |
| --- | --- | --- |
| max_lock_retention_period | string | Specifies the maximum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| min_lock_retention_period | string | Specifies the minimum value that can be set as the retention period for an object in a bucket with locking enabled. The value for this property can be in years or days, not both. The value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years.</num></num> |
| name | string | Specifies the name of the S3 server. A server name can contain 3 to 253 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. Valid values range from 1 to 65535. |

| Name | Type | Description |
|---|---|---|
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. Valid values range from 1 to 65535. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 service bucket configuration

### Manage S3 service bucket configuration

#### Overview

An S3 bucket is a container of objects. Each bucket defines an object namespace. S3 server requests specify objects using a bucket-name and object-name pair. An object consists of data, along with optional metadata and access controls, accessible via a name. An object resides within a bucket. There can be more than one bucket in an S3 server. Buckets which are created for the server are associated with an S3 user that is created on the S3 server.
An access policy is an object that when associated with a resource, defines their permissions. Buckets and objects are defined as resources. By default, only the "root" user can access these resources. Access policies are used to manage access to these resources by enabling ONTAP admin to provide "grants" to allow other users to perform operations on the buckets.

#### Examples

#### Retrieving all fields for all S3 buckets of an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/buckets?fields=**&return_records=true" -H "accept:
application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
    "name": "bucket-2",
    "volume": {
      "name": "fg_oss_1558514455",
      "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
    },
    "size": 107374182400,
    "logical_used_size": 157286400,
    "encryption": {
      "enabled": false
    },
    "comment": "S3 bucket.",
    "qos_policy": {
      "min_throughput_iops": 0,
      "min_throughput_mbps": 0,
      "max_throughput_iops": 1000,
      "max_throughput_mbps": 50,
      "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
      "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
    },
    "audit_event_selector": {
      "access":"all",
      "permission":"all"
    },
    "versioning-state": "disabled",
    "snapshot-policy": {
      "name": "default-1weekly",
      "uuid": "f9c5f090-4ac8-11ef-ba24-005056a7ceb6"
```

```
      },
    },
    {
      "svm": {
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
        "name": "vs1"
      },
      "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
      "name": "bucket-1",
      "volume": {
        "name": "fg_oss_1558690256",
        "uuid": "a36a1ea7-7e06-11e9-97e8-0050568ea123"
      },
      "size": 107374182400,
      "logical_used_size": 0,
      "encryption": {
        "enabled": false
      },
      "comment": "bucket1",
      "qos_policy": {
        "min_throughput_iops": 0,
        "min_throughput_mbps": 0,
        "max_throughput_iops": 1000,
        "max_throughput_mbps": 50,
        "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
        "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
      },
      "policy": {
        "statements": [
          {
            "effect": "allow",
            "actions": [
              "*"
            ],
            "principals": [
              "Alice"
            ],
            "resources": [
              "*"
            ],
            "sid": "fullAccessForAliceToBucket"
          },
          {
            "effect": "allow",
            "actions": [
              "ListBucket",
```

```
            "GetObject"
          ],
          "principals": [
            "ann",
            "jack"
          ],
          "resources": [
            "bucket-1",
            "bucket-1/*"
          ],
          "sid": "AccessToListAndGetObjectForAnnAndJack",
          "conditions": [
            {
              "operator": "ip_address",
              "source_ips": [
                "1.1.1.1/10"
              ]
            },
            {
              "operator": "string_equals",
              "prefixes": [
                "pref1",
                "pref2"
              ],
              "usernames": [
                "user1",
                "user2"
              ],
              "delimiters": [
                "del1",
                "del2"
              ]
            },
            {
              "operator": "numeric_equals",
              "max_keys": [
                100
              ]
            }
          ]
        },
        {
          "effect": "deny",
          "actions": [
            "*Object"
          ],
```

```
          "principals": [
            "mike",
            "group/group1",
            "nasgroup/group2",
          ],
          "resources": [
            "bucket-1/policy-docs/*",
            "bucket-1/confidential-*"
          ],
          "sid": "DenyAccessToGetPutDeleteObjectForMike"
        },
        {
          "effect": "allow",
          "actions": [
            "GetObject"
          ],
          "principals": [
            "*"
          ],
          "resources": [
            "bucket-1/readme"
          ],
          "sid": "AccessToGetObjectForAnonymousUsers"
        },
        {
          "effect": "allow",
          "actions": [
            "GetObject"
          ],
          "principals": [
          ],
          "resources": [
            "bucket-1/policies/examples/*"
          ],
          "sid": "AccessToGetObjectForAllUsersOfSVM"
        }
      ]
    },
    "cors": {
      "rules": [
        {
          "id": "string",
          "allowed_origins": [
            "http://www.example.com"
          ],
          "allowed_methods": [
```

```
            "PUT",
            "DELETE"
          ],
          "allowed_headers": [
            "x-amz-request-id"
          ],
          "expose_headers": [
            "http://www.example.com"
          ],
          "max_age_seconds": 1024
        }
      ]
    }
  }
],
"num_records": 2
}
```


**Retrieving the specified bucket associated with an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/buckets/527812ab-7c6d-11e9-97e8-0050568ea123" -H
"accept: application/json"

# The response:
{
"svm": {
  "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
  "name": "vs1"
},
"uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
"name": "bucket-2",
"volume": {
  "name": "fg_oss_1558514455",
  "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
},
"size": 107374182400,
"logical_used_size": 157286400,
"encryption": {
  "enabled": false
},
"comment": "S3 bucket.",
"qos_policy": {
  "min_throughput_iops": 0,
  "min_throughput_mbps": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
},
"snapshot-policy": {
  "name": "default-1weekly",
  "uuid": "f9c5f090-4ac8-11ef-ba24-005056a7ceb6"
},
"snapshot_restore": {
  "name": "snap1",
  "state": "restoring",
  "progress": 80,
  "objects_remaining": 300
},
}
```

**Creating an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea123\" } ],
\"comment\": \"S3 bucket.\", \"constituents_per_aggregate\": 4, \"name\":
\"bucket-3\"}"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Creating an S3 bucket along with QoS policies and event selector for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:
curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-
0050568e9798/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"comment\": \"S3 bucket.\",
\"name\": \"bucket-3\", \"qos_policy\": { \"min_throughput_iops\": 0,
\"min_throughput_mbps\": 0, \"max_throughput_iops\": 1000000,
\"max_throughput_mbps\": 900000, \"uuid\": \"02d07a93-6177-11ea-b241-
000c293feac8\", \"name\":
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" },
\"audit_event_selector\": { \"access\": \"all\", \"permission\":
\"all\"}}"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

#### Creating an S3 bucket along with policies for an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-
0050568e9798/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea123\" } ],
\"comment\": \"S3 bucket.\", \"constituents_per_aggregate\": 4, \"name\":
\"bucket-3\", \"policy\": { \"statements\": [ { \"actions\": [
\"GetObject\" ], \"conditions\": [ { \"operator\": \"ip_address\",
\"source_ips\": [ \"1.1.1.1/23\", \"1.2.2.2/20\" ] }, { \"max_keys\": [
1000 ], \"operator\": \"numeric_equals\" }, { \"delimiters\": [ \"/\" ],
\"operator\": \"string_equals\", \"prefixes\": [ \"pref\" ],
\"usernames\": [ \"user1\" ] } ], \"effect\": \"allow\", \"resources\": [
\"bucket-3/policies/examples/*\" ], \"sid\":
\"AccessToGetObjectForAllUsersofSVM\" }, { \"actions\": [ \"*Object\" ],
\"effect\": \"deny\", \"principals\": [ \"mike\", \"group/grp1\" ],
\"resources\": [ \"bucket-3/policy-docs/*\", \"bucket-3/confidential-*\"
], \"sid\": \"DenyAccessToObjectForMike\" }, { \"actions\": [
\"GetObject\" ], \"effect\": \"allow\", \"principals\": [ \"*\" ],
\"resources\": [ \"bucket-3/readme\" ], \"sid\":
\"AnonymousAccessToGetObjectForUsers\" } ] } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-3",
    "comment": "S3 bucket."
```

```
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

#### Creating an S3 bucket along with lifecycle management rules

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-
0050568e9798/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea123\" } ],
\"comment\": \"S3 bucket.\", \"constituents_per_aggregate\": 4, \"name\":
\"bucket-4\", \"lifecycle_management\": { \"rules\": [ { \"name\":
\"rule1\", \"expiration\": { \"object_age_days\" : \"1000\" } ,
\"abort_incomplete_multipart_upload\" : { \"after_initiation_days\" : 200
} , \"object_filter\": { \"prefix\" : \"obj1*/\" ,  \"size_greater_than\"
: \"1000\" } }, { \"name\": \"rule2\", \"object_filter\": {
\"size_greater_than\" : \"50\" }, \"expiration\": { \"object_age_days\" :
\"5000\" } } ] } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 18 April 2022 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-4
Content-Length: 363
Content-Type: application/json
```

```
{
"num_records": 1,
"records": [
  {
    "name": "bucket-4",
    "comment": "S3 bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```


#### Creating an S3 bucket with object locking enabled for an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:

curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea143/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea143\" } ],
\"comment\": \"S3 Compliance mode bucket.\",
\"constituents_per_aggregate\": 4, \"name\": \"bucket-5\", \"retention\":
{ \"mode\": \"compliance\", \"default_period\" : \"P1Y\" } }"



# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2023 11:25:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea143/buckets/?name=bucket-5
Content-Length: 353
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "bucket-5",
    "comment": "S3 Compliance mode bucket."
  }
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea143",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea143"
    }
  }
}
}
```

#### Creating an S3 bucket with CORS rules for an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/2da46d66-2327-
11ef-ab73-005056a7a2dd/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H  "authorization: Basic <CREDENTIALS>" -H
"Content-Type: application/json" -d "{  \"cors\": { \"rules\": [{
\"allowed_headers\": [ \"x-amz-request-id\" ], \"allowed_methods\": [
\"PUT\", \"DELETE\" ], \"allowed_origins\": [ \"http://www.example.com\"
], \"expose_headers\": [ \"http://www.example.com\" ], \"id\": \"id1\",
\"max_age_seconds\": 1024 }]}, \"name\": \"bucket1\"}"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 06 Jun 2024 09:18:26 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/protocols/s3/services/2da46d66-2327-11ef-ab73-
005056a7a2dd/buckets/baf69601-23e5-11ef-ab73-005056a7a2dd
Content-Length: 871
Content-Type: application/json
Vary: Origin
{
"num_records": 1,
"records": [
  {
    "uuid": "baf69601-23e5-11ef-ab73-005056a7a2dd",
    "name": "bucket1",
    "type": "s3",
    "cors": {
      "rules": [
        {
          "id": "id1",
          "allowed_origins": [
            "http://www.example.com"
          ],
          "allowed_methods": [
            "PUT",
            "DELETE"
```

```
          ],
          "allowed_headers": [
            "x-amz-request-id"
          ],
          "expose_headers": [
            "http://www.example.com"
          ],
          "max_age_seconds": 1024
        }
      ]
    }
  }
],
"job": {
  "uuid": "baf69953-23e5-11ef-ab73-005056a7a2dd",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/baf69953-23e5-11ef-ab73-005056a7a2dd"
    }
  }
}
}
```


**Creating an S3 bucket with a snapshot-policy**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets

# The call:
curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea143/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"comment\": \"S3 snapshot policy
bucket.\", \"snapshot_policy\": { \"name\": \"default-1weekly\", \"uuid\":
\"f9c5f090-4ac8-11ef-ba24-005056a7ceb6\" }, \"name\": \"bucket-6\" }"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 25 Jul 2024 16:40:37 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/8c38f10b-4871-11ef-aab5-
005056a7ceb6/buckets/9f21f404-4aa4-11ef-ba24-005056a7ceb6
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9f21f716-4aa4-11ef-ba24-005056a7ceb6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9f21f716-4aa4-11ef-ba24-005056a7ceb6"
    }
  }
}
}
```

**Updating an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/buckets/754389d0-7e13-11e9-bfdc-0050568ea122" -H
"accept: application/json?return_records=true" -H "Content-Type:
application/json" -d "{ \"comment\": \"Bucket modified.\", \"size\":
111111111111,  \"qos_policy\": { \"min_throughput_iops\": 0,
\"min_throughput_mbps\": 0, \"max_throughput_iops\": 1000000,
\"max_throughput_mbps\": 900000, \"uuid\": \"02d07a93-6177-11ea-b241-
000c293feac8\", \"name\":
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\"}}"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating an S3 bucket policy and event selector for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-
11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ {
\"actions\": [ \"*\" ], \"conditions\": [ { \"operator\": \"ip_address\",
\"source_ips\": [ \"1.1.1.1/23\", \"1.2.2.2/20\" ] }, { \"max_keys\": [
1000 ], \"operator\": \"numeric_equals\" }, { \"delimiters\": [ \"/\" ],
\"operator\": \"string_equals\", \"prefixes\": [ \"pref\" ],
\"usernames\": [ \"user1\" ] } ], \"effect\": \"allow\", \"resources\": [
\"*\" ], \"sid\": \"fullAccessForAllPrincipalsToBucket\"} ] },
\"audit_event_selector\": { \"access\": \"read\", \"permission\":
\"deny\"} }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating the default-retention period on an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-
11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"retention\": { \"default_period\"
: \"P10Y\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2023 11:35:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea143",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea143"
    }
  }
}
}
```

**Updating the CORS rules of an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/2da46d66-2327-
11ef-ab73-005056a7a2dd/buckets/dce4e40-2330-11ef-ab73-
005056a7a2dd?return_timeout=0" -H  "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H  "Content-Type: application/json"
-d "{ \"cors\": { \"rules\": [{ \"allowed_headers\": [ \"x-amz-new-
header\" ], \"allowed_methods\": [ \"PUT\" ], \"allowed_origins\": [
\"http://www.example.in\" ], \"expose_headers\": [ \"http://www.example\"
], \"id\": \"id3\", \"max_age_seconds\": 10 }]}}"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 06 Jun 2024 09:30:57 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Length: 189
Content-Type: application/json
Vary: Origin
{
"job": {
  "uuid": "7aaa2b64-23e7-11ef-ab73-005056a7a2dd",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/7aaa2b64-23e7-11ef-ab73-005056a7a2dd"
    }
  }
}
}
```

**Updating the snapshot-policy for an S3 bucket for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-
11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"snapshot_policy\": { \"name\":
\"default-1weekly\", \"uuid\": \"f9c5f090-4ac8-11ef-ba24-005056a7ceb6\" }
}"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 25 Jul 2024 16:56:44 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "df350023-4aa6-11ef-ba24-005056a7ceb6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/df350023-4aa6-11ef-ba24-005056a7ceb6"
    }
  }
}
}
```

**Restoring a snapshot for an S3 bucket**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}?restore_to.snapshot.n
ame={snapshot.name}

# The call:
curl -iku admin:<password> -X PATCH "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?restore_to.snapshot.name=snap1" -H "accept: application/json"
-H "Content-Type: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 25 Jul 2024 16:56:44 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "df350023-4aa6-11ef-ba24-005056a7ceb6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/df350023-4aa6-11ef-ba24-005056a7ceb6"
    }
  }
}
}
```

**Deleting an S3 bucket policy for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-
11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ ] }
}"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Deleting CORS rules for an S3 bucket in an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-
11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-
0050568ea122?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"cors\": { \"rules\": [ ] } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2024 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Deleting an S3 bucket for a specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -iku admin:<password> -X DELETE "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/754389d0-7e13-11e9-bfdc-
0050568ea123?return_records=true" -H "accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:40:17 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "b3af4a54-7e18-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b3af4a54-7e18-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

### Retrieve the S3 bucket configurations for an SVM

GET `/protocols/s3/services/{svm.uuid}/buckets`

**Introduced In:** 9.7

Retrieves the S3 bucket's configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket show`
- `vserver object-store-server bucket policy statement show`
- `vserver object-store-server bucket policy-statement-condition show`
- `vserver object-store-server bucket lifecycle-management-rule show`
- `vserver object-store-server bucket snapshot restore show`

**Learn more**

- DOC /protocols/s3/services/{svm.uuid}/buckets

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name<br>• Introduced in: 9.16 |
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid<br>• Introduced in: 9.16 |
| cors.rules.allowed_headers | string | query | False | Filter by cors.rules.allowed_headers<br>• Introduced in: 9.16 |
| cors.rules.id | string | query | False | Filter by cors.rules.id<br>• Introduced in: 9.16<br>• maxLength: 256<br>• minLength: 0 |
| cors.rules.max_age_seconds | integer | query | False | Filter by cors.rules.max_age_seconds<br>• Introduced in: 9.16 |
| cors.rules.allowed_origins | string | query | False | Filter by cors.rules.allowed_origins<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cors.rules.expose_headers | string | query | False | Filter by cors.rules.expose_headers<br>• Introduced in: 9.16 |
| cors.rules.allowed_methods | string | query | False | Filter by cors.rules.allowed_methods<br>• Introduced in: 9.16 |
| snapshot_restore.state | string | query | False | Filter by snapshot_restore.state<br>• Introduced in: 9.18 |
| snapshot_restore.objects_remaining | integer | query | False | Filter by snapshot_restore.objects_remaining<br>• Introduced in: 9.18 |
| snapshot_restore.progress | integer | query | False | Filter by snapshot_restore.progress<br>• Introduced in: 9.18 |
| snapshot_restore.snapshot | string | query | False | Filter by snapshot_restore.snapshot<br>• Introduced in: 9.18 |
| name | string | query | False | Filter by name<br>• maxLength: 63<br>• minLength: 3 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| policy.statements.principals | string | query | False | Filter by policy.statements.principals<br>• Introduced in: 9.8 |
| policy.statements.sid | string | query | False | Filter by policy.statements.sid<br>• Introduced in: 9.8<br>• maxLength: 256<br>• minLength: 0 |
| policy.statements.resources | string | query | False | Filter by policy.statements.resources<br>• Introduced in: 9.8 |
| policy.statements.effect | string | query | False | Filter by policy.statements.effect<br>• Introduced in: 9.8 |
| policy.statements.actions | string | query | False | Filter by policy.statements.actions<br>• Introduced in: 9.8 |
| policy.statements.conditions.operator | string | query | False | Filter by policy.statements.conditions.operator<br>• Introduced in: 9.8 |
| policy.statements.conditions.delimiters | string | query | False | Filter by policy.statements.conditions.delimiters<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| policy.statements.conditions.prefixes | string | query | False | Filter by policy.statements.conditions.prefixes<br>• Introduced in: 9.8 |
| policy.statements.conditions.usernames | string | query | False | Filter by policy.statements.conditions.usernames<br>• Introduced in: 9.8 |
| policy.statements.conditions.source_ips | string | query | False | Filter by policy.statements.conditions.source_ips<br>• Introduced in: 9.8 |
| policy.statements.conditions.max_keys | integer | query | False | Filter by policy.statements.conditions.max_keys<br>• Introduced in: 9.8 |
| lifecycle_management.rules.bucket_name | string | query | False | Filter by lifecycle_management.rules.bucket_name<br>• Introduced in: 9.14<br>• maxLength: 63<br>• minLength: 3 |
| lifecycle_management.rules.non_current_version_expiration.non_current_days | integer | query | False | Filter by lifecycle_management.rules.non_current_version_expiration.non_current_days<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lifecycle_management.rules.non_current_version_expiration.new_non_current_versions | integer | query | False | Filter by lifecycle_management.rules.non_current_version_expiration.new_non_current_versions<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.uuid | string | query | False | Filter by lifecycle_management.rules.uuid<br><br>• Introduced in: 9.14 |
| lifecycle_management.rules.expiration.expired_object_delete_marker | boolean | query | False | Filter by lifecycle_management.rules.expiration.expired_object_delete_marker<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.expiration.object_expiry_date | string | query | False | Filter by lifecycle_management.rules.expiration.object_expiry_date<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.expiration.object_age_days | integer | query | False | Filter by lifecycle_management.rules.expiration.object_age_days<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lifecycle_management.rules.name | string | query | False | Filter by lifecycle_management.rules.name<br>• Introduced in: 9.13<br>• maxLength: 256<br>• minLength: 0 |
| lifecycle_management.rules.svm.name | string | query | False | Filter by lifecycle_management.rules.svm.name<br>• Introduced in: 9.14 |
| lifecycle_management.rules.svm.uuid | string | query | False | Filter by lifecycle_management.rules.svm.uuid<br>• Introduced in: 9.14 |
| lifecycle_management.rules.object_filter.prefix | string | query | False | Filter by lifecycle_management.rules.object_filter.prefix<br>• Introduced in: 9.13 |
| lifecycle_management.rules.object_filter.size_greater_than | integer | query | False | Filter by lifecycle_management.rules.object_filter.size_greater_than<br>• Introduced in: 9.13 |
| lifecycle_management.rules.object_filter.size_less_than | integer | query | False | Filter by lifecycle_management.rules.object_filter.size_less_than<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lifecycle_management.rules.object_filter.tags | string | query | False | Filter by lifecycle_management.rules.object_filter.tags<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.enabled | boolean | query | False | Filter by lifecycle_management.rules.enabled<br><br>• Introduced in: 9.13 |
| lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days | integer | query | False | Filter by lifecycle_management.rules.abort_incomplete_multipart_upload.after_initiation_days<br><br>• Introduced in: 9.13 |
| audit_event_selector.permission | string | query | False | Filter by audit_event_selector.permission<br><br>• Introduced in: 9.10 |
| audit_event_selector.access | string | query | False | Filter by audit_event_selector.access<br><br>• Introduced in: 9.10 |
| type | string | query | False | Filter by type<br><br>• Introduced in: 9.12 |
| versioning_state | string | query | False | Filter by versioning_state<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| is_consistent_etag | boolean | query | False | Filter by is_consistent_etag<br>• Introduced in: 9.17 |
| qos_policy.min_throughput_mbps | integer | query | False | Filter by qos_policy.min_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| qos_policy.max_throughput_iops | integer | query | False | Filter by qos_policy.max_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| qos_policy.max_throughput_mbps | integer | query | False | Filter by qos_policy.max_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| qos_policy.name | string | query | False | Filter by qos_policy.name<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| qos_policy.min_throughput_iops | integer | query | False | Filter by qos_policy.min_throughput_iops<br>• Introduced in: 9.8<br>• Max value: 2147483647<br>• Min value: 0 |
| qos_policy.uuid | string | query | False | Filter by qos_policy.uuid<br>• Introduced in: 9.8 |
| qos_policy.max_throughput | string | query | False | Filter by qos_policy.max_throughput<br>• Introduced in: 9.17 |
| qos_policy.min_throughput | string | query | False | Filter by qos_policy.min_throughput<br>• Introduced in: 9.17 |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |
| logical_used_size | integer | query | False | Filter by logical_used_size |
| uuid | string | query | False | Filter by uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| size | integer | query | False | Filter by size<br>• Max value: 62672162783232000<br>• Min value: 199229440 |
| role | string | query | False | Filter by role<br>• Introduced in: 9.10 |
| retention.default_period | string | query | False | Filter by retention.default_period<br>• Introduced in: 9.14 |
| retention.mode | string | query | False | Filter by retention.mode<br>• Introduced in: 9.14 |
| svm.name | string | query | False | Filter by svm.name |
| protection_status.destination.is_cloud | boolean | query | False | Filter by protection_status.destination.is_cloud<br>• Introduced in: 9.10 |
| protection_status.destination.is_ontap | boolean | query | False | Filter by protection_status.destination.is_ontap<br>• Introduced in: 9.10 |
| protection_status.destination.is_external_cloud | boolean | query | False | Filter by protection_status.destination.is_external_cloud<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protection_status.is_protected | boolean | query | False | Filter by protection_status.is_protected<br><br>• Introduced in: 9.10 |
| encryption.enabled | boolean | query | False | Filter by encryption.enabled |
| nas_path | string | query | False | Filter by nas_path<br><br>• Introduced in: 9.12 |
| is_nas_path_mutable | boolean | query | False | Filter by is_nas_path_mutable<br><br>• Introduced in: 9.17 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_bucket_svm] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "audit_event_selector": {
        "access": "string",
        "permission": "string"
      },
      "comment": "S3 bucket.",
      "cors": {
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "allowed_headers": [
              "x-amz-request-id"
            ],
            "allowed_methods": [
              "PUT",
              "DELETE"
            ],
            "allowed_origins": [
              "http://www.example.com"
            ],
            "expose_headers": [
              "x-amz-date"
            ],
            "id": "string",
            "max_age_seconds": 1024
          }
        ]
      },
      "lifecycle_management": {
```

```
        "rules": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "abort_incomplete_multipart_upload": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              }
            },
            "bucket_name": "bucket1",
            "expiration": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "object_age_days": 100,
              "object_expiry_date": "2039-09-23 00:00:00 +0000"
            },
            "name": "string",
            "non_current_version_expiration": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              }
            },
            "object_filter": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "prefix": "/logs",
              "size_greater_than": 10240,
              "size_less_than": 10485760,
              "tags": [
                "project1=projA",
                "project2=projB"
              ]
            },
```

```
            "svm": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "svm1",
              "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
            },
            "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
          }
        ]
      },
      "logical_used_size": 0,
      "name": "bucket1",
      "nas_path": "/",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  1000
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ],
```

```
            "effect": "allow",
            "principals": [
              "user1",
              "group/grp1",
              "nasgroup/group1"
            ],
            "resources": [
              "bucket1",
              "bucket1/*"
            ],
            "sid": "Full_Access_To_User1!"
          }
        ]
      },
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "retention": {
```

```
          "default_period": "P10Y",
          "mode": "governance"
        },
        "role": "string",
        "size": 819200000,
        "snapshot_policy": {
          "name": "default-1weekly",
          "uuid": "3675af31-431c-12fa-114a-20675afebc12"
        },
        "snapshot_restore": {
          "snapshot": "string",
          "state": "string"
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "type": "s3",
        "uuid": "string",
        "versioning_state": "enabled",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be set in a POST or PATCH method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |

| Name | Type | Description |
| --- | --- | --- |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. This field cannot be set using the PATCH method. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
|---|---|---|
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
|---|---|---|
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP.<br>• Default value:<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be set in a POST method.<br>• Default value:<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket_svm

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |

| Name | Type | Description |
| --- | --- | --- |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |

| Name | Type | Description |
|---|---|---|
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an S3 bucket configuration for an SVM

POST `/protocols/s3/services/{svm.uuid}/buckets`

**Introduced In:** 9.7

Creates the S3 bucket configuration of an SVM.

#### Important notes

- Each SVM can have one or more bucket configurations.
- Aggregate lists should be specified explicitly. If not specified, then the bucket is auto-provisioned as a FlexGroup.
- Constituents per aggregate specifies the number of components (or FlexVols) per aggregate. Is specified only when an aggregate list is explicitly defined.
- An access policy can be created when a bucket is created.
- "qos_policy" can be specified if a bucket needs to be attached to a QoS group policy during creation time.
- "audit_event_selector" can be specified if a bucket needs to be specify access and permission type for auditing.
- Cross-origin resource sharing (CORS) configuration can be specified when a bucket is created.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the bucket configuration.
- `name` - Bucket name that is to be created.

#### Recommended optional properties

- `aggregates` - List of aggregates for the FlexGroup on which the bucket is hosted on.
- `constituents_per_aggregate` - Number of constituents per aggregate.
- `size` - Specifying the bucket size is recommended.
- `policy` - Specifying policy enables users to perform operations on buckets; specifying the resource permissions is recommended.
- `qos_policy` - A QoS policy for buckets.
- `audit_event_selector` - Audit policy for buckets.

- `versioning_state` - Versioning state for buckets.
- `type` - Type of bucket.
- `nas_path` - The NAS path to which the NAS bucket corresponds to.
- `use_mirrored_aggregates` - Specifies whether mirrored aggregates are selected when provisioning a FlexGroup volume.
- `lifecycle_management` - Object store server lifecycle management policy.
- `retention.mode` - Object lock mode supported on the bucket.
- `retention.default_period` - Specifies the duration of default-retention applicable for objects on the object store bucket.
- `cors` - Specifying CORS rules enables the bucket to service the cross-origin requests.
- `snapshot_policy` - Snapshot policy for the bucket.
- `is_consistent_etag` - Return a consistent ETag for NAS buckets.
- `is_nas_path_mutable` - Specifies whether the NAS bucket mapping with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount.
- `snapshot_restore` - Specifies information regarding a snapshot restore operation on the bucket.

**Default property values**

- `size` - 800GB
- `comment` - ""
- `aggregates` - No default value.
- `constituents_per_aggregate` - *4* , if an aggregates list is specified. Otherwise, no default value.
- `policy.statements.actions` - GetObject, PutObject, DeleteObject, ListBucket, ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging, PutObjectTagging, DeleteObjectTagging, GetBucketVersioning, PutBucketVersioning.
- `policy.statements.principals` - all S3 users and groups in the SVM or the NAS groups.
- `policy.statements.resources` - all objects in the bucket.
- `policy.statements.conditions` - list of bucket policy conditions.
- `qos-policy` - No default value.
- `versioning_state` - disabled.
- `use_mirrored_aggregates` - *true* for a MetroCluster configuration and *false* for a non-MetroCluster configuration.
- `type` - S3.
- `retention.mode` - no_lock
- `is_consistent_etag` - false

#### Related ONTAP commands

- `vserver object-store-server bucket create`
- `vserver object-store-server bucket policy statement create`
- `vserver object-store-server bucket policy-statement-condition create`
- `vserver object-store-server bucket lifecycle-management-rule create`
- `vserver object-store-server bucket cors-rule create`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| aggregates | array[aggregates] | A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume. The "uuid" field cannot be used with the field "storage_service_level". |
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| constituents_per_aggregate | integer | Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently. This field cannot be set using the PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |

| Name | Type | Description |
| --- | --- | --- |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| storage_service_level | string | Specifies the storage service level of the FlexGroup volume on which the bucket should be created. Valid values are "value", "performance" or "extreme". This field cannot be used with the field "aggregates.uuid" or with the "constituents_per_aggregate" in a POST method. This field cannot be set using the PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

**Example request**

```
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "constituents_per_aggregate": 4,
  "cors": {
    "rules": [
      {
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "lifecycle_management": {
    "rules": [
      {
        "bucket_name": "bucket1",
        "expiration": {
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "name": "string",
        "object_filter": {
```

```
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
    ]
  },
  "logical_used_size": 0,
  "name": "bucket1",
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
```

```
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
      }
    ]
  },
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y",
    "mode": "governance"
```

```
  },
  "role": "string",
  "size": 819200000,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
    "state": "string"
  },
  "storage_service_level": "value",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "s3",
  "uuid": "string",
  "versioning_state": "enabled",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**


```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405777 + | "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason of failure}. "; |
| 92405785 + | "Bucket name "{bucket name}" contains invalid characters or invalid character combinations. Valid characters for a bucket name are 0-9, a-z, ".", and "-". Invalid character combinations are ".-", "-.", and "..". "; |
| 92405786 + | "Bucket name "{bucket name}" is not valid. Bucket names must have between 3 and 63 characters. "; |
| 92405811 + | "Failed to create bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |
| 92405812 + | "Failed to create the object store volume. Reason: {Reason for failure}."; |
| 92405819 + | "Cannot provision an object store server volume for bucket "{bucket name}" in SVM "{svm.name}" on the following aggregates because they are SnapLock aggregates: {List of aggregates.name}."; |
| 92405820 + | "Failed to check whether the aggregate "{aggregates.name}" is a FabricPool. Reason: {Reason for failure}."; |
| 92405821 + | "Cannot provision an object store server volume for bucket "{bucket name}" in SVM "{svm.name}" on the following aggregates because they are FabricPool: {List of aggregates.name}."; |
| 92405827 + | "Internal Error. Unable to generate object store volume name."; |

| Error code | Message |
| --- | --- |
| 92405857 + | "One or more aggregates must be specified if "constituents_per_aggregate" is specified."; |
| 92405858 + | "Failed to "create" the "bucket" because the operation is only supported on data SVMs."; |
| 92405859 + | "The specified "aggregates.uuid" "{aggregates.uuid}" does not exist."; |
| 92405860 + | "The specified "aggregates.name" "{aggregates.name}" and "aggregates.uuid" "{aggregates.uuid}" do not refer to same aggregate."; |
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92405863 + | "Failed to create access policies for bucket "{bucket name}". Reason: {Reason of Failure}. Resolve all the issues and retry the operation."; |
| 92405863 + | "Failed to create event selector for bucket "{bucket name}". Reason: {Reason of failure}. Reason: object-store-server audit configuration not created for SVM "{svm.name}". Resolve all the issues and retry the operation."; |
| 92405891 + | The resources specified in the access policy are not valid. Valid ways to specify a resource are *, <bucket-name>, <bucket-name>/…/…. Valid characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &, @,?, (, ), single quote, *, !, - and $. |
| 92406166 + | "Cannot enable locking on a NAS bucket."; |
| 92406170 + | "Cannot set "-default-retention-period" on object store bucket "{0}" in Vserver "{1}". Setting the default retention period on an object store bucket requires an effective cluster version of 9.14.1 or later."; |
| 92406171 + | "Cannot set "{retention_mode}" to "compliance" in a MetroCluster configuration"; |
| 92406174 + | "Internal error. Failed to complete bucket create workflow with "-retention-mode" set to "compliance" or "governance". Reason: {0}"; |
| 92406175 + | "The SnapLock compliance clock is not running. Use the "snaplock compliance-clock initialize" command to initialize the compliance clock, and then try the operation again."; |
| 92406176 + | "The SnapLock compliance clock is not running on the MetroCluster partner cluster. Use the "snaplock compliance-clock initialize" command to initialize the compliance clock on the MetroCluster partner cluster, and then try the operation again."; |

| Error code | Message |
| --- | --- |
| 92406230 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be greater than the maximum lock retention period set in the object store server for SVM "{SVM}". Check the maximum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406236 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be less than the minimum lock retention period set in the object store server for SVM "{SVM}". Check the minimum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406217 + | "The specified "allowed_headers" is not valid because it contains more than one wild card ("*") character."; |
| 92406224 + | "A Cross-Origin Resource Sharing (CORS) rule must have an origin and HTTP method specified."; |
| 92406222 + | "Cannot specify Cross-Origin Resource Sharing (CORS) configuration for object store bucket "{bucket}" on SVM "{SVM}". Specifying such configuration is supported on object store volumes created in ONTAP 9.8 or later releases only."; |
| 92406211 + | "The specified method "DONE" is not valid. Valid methods are GET, PUT, DELETE, HEAD, and POST."; |
| 92405863 + | "Failed to create CORS rules for bucket "bb1". Reason: "Field "index" cannot be specified for this operation.". Resolve all the issues and retry the operation."; |
| 92406228 + | "Cannot exceed the maximum limit of 100 Cross-Origin Resource Sharing (CORS) rules per S3 bucket "{bucket}" in SVM "{SVM}".";; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be set in a POST or PATCH method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. This field cannot be set using the PATCH method. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |

| Name | Type | Description |
| --- | --- | --- |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be set in a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or

"qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
|---|---|---|
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket_svm

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume. The "uuid" field cannot be used with the field "storage_service_level". |
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| constituents_per_aggregate | integer | Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently. This field cannot be set using the PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |

| Name | Type | Description |
| --- | --- | --- |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |

| Name | Type | Description |
| --- | --- | --- |
| storage_service_level | string | Specifies the storage service level of the FlexGroup volume on which the bucket should be created. Valid values are "value", "performance" or "extreme". This field cannot be used with the field "aggregates.uuid" or with the "constituents_per_aggregate" in a POST method. This field cannot be set using the PATCH method. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 bucket configuration for an SVM

DELETE `/protocols/s3/services/{svm.uuid}/buckets/{uuid}`

**Introduced In:** 9.7

Deletes the S3 bucket configuration of an SVM. An access policy is also deleted on an S3 bucket "delete" command. All CORS rules are also deleted on bucket deletion.

#### Related ONTAP commands

- `vserver object-store-server bucket delete`
- `vserver object-store-server bucket policy statement delete`
- `vserver object-store-server bucket policy-statement-condition delete`
- `vserver object-store-server bucket lifecycle-management-rule delete`
- `vserver object-store-server bucket cors-rule delete`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405811 + | "Failed to delete bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |
| 92405858 + | "Failed to "delete" the "bucket" because the operation is only supported on data SVMs."; |
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92405779 + | "Failed to remove bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason for failure}. "; |
| 92405813 + | "Failed to delete the object store volume. Reason: {Reason for failure}."; |
| 92405864 + | "An error occurred when deleting an access policy. The reason for failure is detailed in the error message."; |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve an S3 bucket for an SVM**

GET `/protocols/s3/services/{svm.uuid}/buckets/{uuid}`

**Introduced In:** 9.7

Retrieves the S3 bucket configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket show`
- `vserver object-store-server bucket policy statement show`
- `vserver object-store-server bucket policy-statement-condition show`
- `vserver object-store-server bucket lifecycle-management-rule show`
- `vserver object-store-server bucket cors-rule show`
- `vserver object-store-server bucket snapshot restore show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| is_nas_path_mutable | boolean | Specifies whether the NAS bucket mapping or association with a NAS volume can change according to the changes in the NAS volume junction-path due to volume operations like mount and unmount and therefore the NAS bucket will have access to any path in a NAS volume that matches the specified nas-path. Or is immutable and therefore the NAS bucket will always have access to the same nas-path that was specified during bucket creation even if the volume junction-path has undergone changes after the bucket creation. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |

| Name | Type | Description |
| --- | --- | --- |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

**Example response**

```
{
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "cors": {
    "rules": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "lifecycle_management": {
    "rules": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "abort_incomplete_multipart_upload": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
            }
          }
        },
        "bucket_name": "bucket1",
        "expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "name": "string",
        "non_current_version_expiration": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        },
        "object_filter": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
```

```
    ]
  },
  "logical_used_size": 0,
  "name": "bucket1",
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
```

```
      }
    ]
  },
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y",
    "mode": "governance"
  },
  "role": "string",
  "size": 819200000,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
```

```
      "state": "string"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "type": "s3",
    "uuid": "string",
    "versioning_state": "enabled",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

audit_event_selector

Event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |

| Name | Type | Description |
| --- | --- | --- |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be set in a POST or PATCH method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| prefix | string | A prefix that is matched against object-names within a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. This field cannot be set using the PATCH method. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be set in a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |
| mode | string | The lock mode of the bucket.<br> compliance ‐ A SnapLock Compliance (SLC) bucket provides the highest level of WORM protection and an administrator cannot destroy a compliance bucket if it contains unexpired WORM objects.<br> governance ‐ An administrator can delete a Governance bucket.<br> no_lock ‐ Indicates the bucket does not support object locking. |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
|---|---|---|
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an S3 bucket configuration for an SVM

PATCH `/protocols/s3/services/{svm.uuid}/buckets/{uuid}`

**Introduced In:** 9.7

Updates the S3 bucket configuration of an SVM.

#### Important notes

- The following fields can be modified for a bucket:
  - `comment` - Any information related to the bucket.
  - `size` - Bucket size.

- `policy` - An access policy for resources (buckets and objects) that defines their permissions. New policies are created after existing policies are deleted. To retain any of the existing policy statements, you need to specify those statements again. Policy conditions can also be modified using this API.
- `qos_policy` - A QoS policy for buckets.
- `audit_event_selector` - Audit policy for buckets. None can be specified for both access and permission to remove audit event selector.
- `versioning_state` - Versioning state for buckets.
- `nas_path` - NAS path to which the NAS bucket corresponds to.
- `retention.default_period` - Specifies the duration of default-retention applicable for objects on the object store bucket.
- `cors` - Specifying CORS rules enables the bucket to service the cross-origin requests. New CORS rules are created after existing rules are deleted. To retain any of the existing rules, you need to specify those CORS rules again. To remove all the existing CORS rules, specify an empty CORS rules list.
- `snapshot_policy` - Snapshot policy for the bucket.
- `is_consistent_etag` - Return a consistent ETag for NAS buckets.
- `restore_to` - Start a snapshot restore operation for S3 bucket

**Related ONTAP commands**

- `vserver object-store-server bucket modify`
- `vserver object-store-server bucket policy statement modify`
- `vserver object-store-server bucket policy-statement-condition modify`
- `vserver object-store-server bucket cors-rule create`
- `vserver object-store-server bucket cors-rule delete`
- `vserver object-store-server bucket snapshot restore start`

**Learn more**

- DOC /protocols/s3/services/{svm.uuid}/buckets

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the bucket. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| restore_to.snapshot.name | string | query | False | Name of the snapshot to restore. The bucket will be restored to the point in time the snapshot was taken.<br><br>• Introduced in: 9.18 |
| restore_to.snapshot.uuid | string | query | False | UUID of the snapshot to restore. The bucket will be restored to the point in time the snapshot was taken.<br><br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |

| Name | Type | Description |
| --- | --- | --- |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

**Example request**

```
{
  "audit_event_selector": {
    "access": "string",
    "permission": "string"
  },
  "comment": "S3 bucket.",
  "cors": {
    "rules": [
      {
        "allowed_headers": [
          "x-amz-request-id"
        ],
        "allowed_methods": [
          "PUT",
          "DELETE"
        ],
        "allowed_origins": [
          "http://www.example.com"
        ],
        "expose_headers": [
          "x-amz-date"
        ],
        "id": "string",
        "max_age_seconds": 1024
      }
    ]
  },
  "is_nas_path_mutable": true,
  "lifecycle_management": {
    "rules": [
      {
        "expiration": {
          "object_age_days": 100,
          "object_expiry_date": "2039-09-23 00:00:00 +0000"
        },
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
    ]
  },
  "logical_used_size": 0,
```

```
  "nas_path": "/",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              1000
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1",
          "nasgroup/group1"
        ],
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "Full_Access_To_User1!"
      }
    ]
  },
  "qos_policy": {
```

```
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "retention": {
    "default_period": "P10Y"
  },
  "role": "string",
  "size": 819200000,
  "snapshot_policy": {
    "name": "default-1weekly",
    "uuid": "3675af31-431c-12fa-114a-20675afebc12"
  },
  "snapshot_restore": {
    "snapshot": "string",
    "state": "string"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "s3",
  "uuid": "string",
  "versioning_state": "enabled",
  "volume": {
```

```
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405778 + | "Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Reason: {Reason for failure}. "; |
| 92405846 + | "Failed to modify the object store volume. Reason: {Reason for failure}."; |
| 92405811 + | "Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Wait a few minutes and try the operation again."; |

| Error code | Message |
| --- | --- |
| 92405858 + | "Failed to "modify" the "bucket" because the operation is only supported on data SVMs."; |
| 92405861 + | "The specified "{parameter field}", "{parameter value}", does not exist."; |
| 92405863 + | "Failed to create access policies for bucket "{bucket name}". Reason: "Principal list can have a maximum of 10 entries.". Resolve all the issues and retry the operation."; |
| 92405864 + | "An error occurs when deleting an access policy. The reason for failure is detailed in the error message."; |
| 92405891 + | The resources specified in the access policy are not valid. Valid ways to specify a resource are *, <bucket-name>, <bucket-name>/…/…. Valid characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &, @,?, (, ), single quote, *, !, - and $. |
| 92405894 + | "Statements, principals and resources list can have a maximum of 10 entries."; |
| 92405897 + | The principals specified in the access policy are not in the correct format. User name must be in between 1 and 64 characters. Valid characters for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - . |
| 92405898 + | "The SID specified in the access policy is not valid."; |
| 92405940 + | "The specified condition key is not valid for operator "ip-address". Valid choices of keys for this operator: source-ips."; |
| 92406014 + | "Failed to modify event selector for bucket "{bucket name}". If value of either access or permission is set to none, then the other must be set to none as well."; |
| 92406032 + | "Modifying the NAS path for a NAS bucket is not supported."; |
| 92406230 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be greater than the maximum lock retention period set in the object store server for SVM "{SVM}". Check the maximum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406236 + | "The value for "retention.default_period" parameter for object store bucket "{bucket}" cannot be less than the minimum lock retention period set in the object store server for SVM "{SVM}". Check the minimum allowed lock retention period present in the object store server for SVM "{SVM}" and try the operation again."; |
| 92406217 + | "The specified "allowed_headers" is not valid because it contains more than one wild card ("*") character."; |

| Error code | Message |
| --- | --- |
| 92406224 + | "A Cross-Origin Resource Sharing (CORS) rule must have an origin and HTTP method specified."; |
| 92406222 + | "Cannot specify Cross-Origin Resource Sharing (CORS) configuration for object store bucket "{bucket}" on SVM "{SVM}". Specifying such configuration is supported on object store volumes created in ONTAP 9.8 or later releases only."; |
| 92406211 + | "The specified method "DONE" is not valid. Valid methods are GET, PUT, DELETE, HEAD, and POST."; |
| 92405863 + | "Failed to create CORS rules for bucket "bb1". Reason: "Field "index" cannot be specified for this operation.". Resolve all the issues and retry the operation."; |
| 92406228 + | "Cannot exceed the maximum limit of 100 Cross-Origin Resource Sharing (CORS) rules per S3 bucket "{bucket}" in SVM "{SVM}".";; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

audit_event_selector

Event selector allows you to specify access and permission types to audit.

| Name | Type | Description |
|---|---|---|
| access | string | Specifies read and write access types. |
| permission | string | Specifies allow and deny permission types. |

rules

Information about the CORS rule of an S3 bucket.

| Name | Type | Description |
|---|---|---|
| allowed_headers | array[string] | An array of HTTP headers allowed in the cross-origin requests. |
| allowed_methods | array[string] | An array of HTTP methods allowed in the cross-origin requests. |
| allowed_origins | array[string] | List of origins from where a cross-origin request is allowed to originate from for the S3 bucket. |

| Name | Type | Description |
| --- | --- | --- |
| expose_headers | array[string] | List of extra headers sent in the response that customers can access from their applications. |
| id | string | Bucket CORS rule identifier. The length of the name can range from 0 to 256 characters. |
| max_age_seconds | integer | The time in seconds for your browser to cache the preflight response for the specified resource. |

cors

Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store bucket CORS rule. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket. This field cannot be set in a POST or PATCH method. |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

rules

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

lifecycle_management

Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects.

| Name | Type | Description |
| --- | --- | --- |
| rules | array[rules] | Specifies an object store lifecycle management policy. This field cannot be set using the PATCH method. |

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

| Name | Type | Description |
| --- | --- | --- |
| delimiters | array[string] | An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified. |
| max_keys | array[integer] | An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified. |
| operator | string | Condition operator that is applied to the specified condition key. |
| prefixes | array[string] | An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified. |
| source_ips | array[string] | An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified. |
| usernames | array[string] | An array of usernames that a current user in the context is evaluated against using the condition operators. |

s3_bucket_policy_statement

Specifies information about a single access permission.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | |
| conditions | array[s3_bucket_policy_condition] | Specifies bucket policy conditions. |
| effect | string | Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| principals | array[string] | |
| resources | array[string] | |
| sid | string | Specifies the statement identifier used to differentiate between statements. The SID length can range from 1 to 256 characters. |

policy

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

| Name | Type | Description |
| --- | --- | --- |
| statements | array[s3_bucket_policy_statement] | Specifies bucket access policy statement. |

destination

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a bucket is protected within the Cloud. |
| is_external_cloud | boolean | Specifies whether a bucket is protected on external Cloud providers. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a bucket is protected within ONTAP.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

protection_status

Specifies attributes of bucket protection.

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | |
| is_protected | boolean | Specifies whether a bucket is a source and if it is protected within ONTAP and/or an external cloud. This field cannot be set in a POST method.<br>• Default value: 1<br>• readOnly: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

qos_policy

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

retention

Specifies the retention mode and default retention period configured on the bucket.

| Name | Type | Description |
| --- | --- | --- |
| default_period | string | Specifies the default retention period that is applied to objects while committing them to the WORM state without an associated retention period. The retention period can be in years, or days. The retention period value represents a duration and must be specified in the ISO-8601 duration format. A period specified for years and days is represented in the ISO-8601 format as "P<num>Y" and "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. The period string must contain only a single time element that is, either years, or days. A duration which combines different periods is not supported, for example "P1Y10D" is not supported.</num></num> |

snapshot_policy

Specifies the bucket snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the snapshot policy. |
| uuid | string | Specifies the unique identifier of the snapshot policy. |

snapshot_restore

Specifies information regarding a snapshot restore operation on the bucket

| Name | Type | Description |
| --- | --- | --- |
| objects_remaining | integer | Remaining objects to be restored for the bucket |
| progress | integer | Snapshot restore progress in percent |
| snapshot | string | Name of the snapshot being restored for the bucket |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Snapshot restore state of the bucket |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

s3_bucket_svm

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| audit_event_selector | audit_event_selector | Event selector allows you to specify access and permission types to audit. |
| comment | string | Can contain any additional information about the bucket being created or modified. |
| cors | cors | Cross-origin resource sharing (CORS) specifies an object associated with a bucket. The CORS configuration enables the bucket to service the cross-origin requests. A request might typically come from an origin with a domain that is different to that of the bucket. By configuring a CORS rule, you can define a combination of allowed origins, HTTP headers and methods that a bucket can use to filter out the cross-origin requests that it can service successfully. |
| encryption | encryption | |
| is_consistent_etag | boolean | Specifies whether the NAS bucket returns a consistent ETag across different S3 requests. |
| lifecycle_management | lifecycle_management | Lifecycle management is implemented as an object associated with a bucket. It defines rules to be applied against objects within a bucket. These rules are applied in the background and can delete objects. |
| logical_used_size | integer | Specifies the bucket logical used size up to this point. This field cannot be set using the PATCH method. |
| nas_path | string | Specifies the NAS path to which the nas bucket corresponds to. |

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied. |
| protection_status | protection_status | Specifies attributes of bucket protection. |
| qos_policy | qos_policy | Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps". Specifying "min_throughput_iops" or "min_throughput_mbps" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached. |
| retention | retention | Specifies the retention mode and default retention period configured on the bucket. |
| role | string | Specifies the role of the bucket. This field cannot be set in a POST method. |
| size | integer | Specifies the bucket size in bytes; ranges from 190MB to 62PB. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | Specifies the bucket snapshot policy. |
| snapshot_restore | snapshot_restore | Specifies information regarding a snapshot restore operation on the bucket |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Specifies the bucket type. Valid values are "s3"and "nas". This field cannot be set using the PATCH method. |
| uuid | string | Specifies the unique identifier of the bucket. This field cannot be specified in a POST or PATCH method. |
| versioning_state | string | Specifies the versioning state of the bucket. Valid values are "disabled", "enabled" or "suspended". Note that the versioning state cannot be modified to 'disabled' from any other state. |
| volume | volume | Specifies the FlexGroup volume name and UUID where the bucket is hosted. This field cannot be set in a POST method. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 bucket lifecycle management rules

### Manage S3 bucket lifecycle management rules

#### Overview

An S3 bucket lifecycle management rule is a list of objects. Each rule defines a set of actions to be performed on the object within the bucket.

### Adding a lifecycle management rule on a S3 bucket under an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}/rules

# The call:
curl -iku admin:<password> -X POST "https://<mgmt-
ip>/api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules?return_timeout=0&return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"name\":
\"rule1\", \"expiration\" : { \"object_age_days\" : \"1000\" },
\"abort_incomplete_multipart_upload\" : { \"after_initiation_days\" : 200
} , \"object_filter\": { \"prefix\" : \"obj1*/\" ,  \"size_greater_than\"
: \"1000\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 30 November 2022 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules?name=bucket-4
Content-Length: 353
Content-Type: application/json
{
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
```

**Updating a lifecycle management rule on a S3 bucket under an SVM**

```
# The call:
curl -iku admin:<password> -X PATCH "https://<mgmt-
ip>/api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules/rule1?return_timeout=0&return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"expiration\"
: { \"object_age_days\" : \"3000\" },
\"abort_incomplete_multipart_upload\" : { \"after_initiation_days\" :
\"5000\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 30 November 2022 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules?name=bucket-4
Content-Length: 353
Content-Type: application/json
{
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea124",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea124"
    }
  }
}
}
```

**Deleting a lifecycle management rule on a S3 bucket under an SVM**

```
# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/259b4e78-2893-
67ea-9785-890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules/rule1"  -H "accept: application/json"
```

**To delete an action within a rule, pass null inside the action-object.**

```
# The call:
curl -iku admin:<password> -X PATCH "https://<mgmt-
ip>/api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules/rule1?return_timeout=0&return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"expiration\"
: null }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 30 November 2022 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/259b4e78-2893-67ea-9785-
890456bbbec4/buckets/259b4e46-2893-67ea-9145-
909456bbbec4/rules?name=bucket-4
Content-Length: 353
Content-Type: application/json
{
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea124",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea124"
    }
  }
}
}
```

### Retrieve S3 lifecycle rules for a bucket

GET `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules`

**Introduced In:** 9.13

Retrieves all S3 Lifecycle rules associated with a bucket. Note that in order to retrieve S3 bucket rule parameters, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket lifecycle-management-rule show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| bucket_name | string | query | False | Filter by bucket_name<br>• maxLength: 63<br>• minLength: 3<br>• Introduced in: 9.14 |
| non_current_version_expiration.non_current_days | integer | query | False | Filter by non_current_version_expiration.non_current_days |
| non_current_version_expiration.new_non_current_versions | integer | query | False | Filter by non_current_version_expiration.new_non_current_versions |
| uuid | string | query | False | Filter by uuid<br>• Introduced in: 9.14 |
| expiration.expired_object_delete_marker | boolean | query | False | Filter by expiration.expired_object_delete_marker |
| expiration.object_expiry_date | string | query | False | Filter by expiration.object_expiry_date |
| expiration.object_age_days | integer | query | False | Filter by expiration.object_age_days |
| name | string | query | False | Filter by name<br>• maxLength: 256<br>• minLength: 0 |
| svm.name | string | query | False | Filter by svm.name<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| object_filter.prefix | string | query | False | Filter by object_filter.prefix |
| object_filter.size_greater_than | integer | query | False | Filter by object_filter.size_greater_than |
| object_filter.size_less_than | integer | query | False | Filter by object_filter.size_less_than |
| object_filter.tags | string | query | False | Filter by object_filter.tags |
| enabled | boolean | query | False | Filter by enabled |
| abort_incomplete_multipart_upload.after_initiation_days | integer | query | False | Filter by abort_incomplete_multipart_upload.after_initiation_days |
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_bucket_lifecycle_rule] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "abort_incomplete_multipart_upload": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "bucket_name": "bucket1",
      "expiration": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "object_age_days": 100,
        "object_expiry_date": "2039-09-23 00:00:00 +0000"
      },
      "name": "string",
      "non_current_version_expiration": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "object_filter": {
        "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "prefix": "/logs",
          "size_greater_than": 10240,
          "size_less_than": 10485760,
          "tags": [
            "project1=projA",
            "project2=projB"
          ]
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_bucket_lifecycle_rule

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Specifies the unique identifier of the bucket. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create the S3 bucket lifecycle rule configuration

POST `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules`

**Introduced In:** 9.13

Creates the S3 bucket lifecycle rule configuration.

#### Required properties

- `name` - Lifecycle Management rule to be created.
- `expiration` and/or `non_current_version_expiration` and/or `abort_incomplete_multipart_upload`
  - Lifecycle Management actions associated with the lifecycle rule.

#### Recommended optional properties

- `enabled` - Lifecycle Management rule is enabled or not.
- `object_filter.prefix` - Lifecycle Management rule filter prefix.
- `object_filter.tags` - Lifecycle Management rule filter tags.

- `object_filter.size_greater_than` - Lifecycle Management rule filter minimum object size.
- `object_filter.size_less_than` - Lifecycle Management rule filter maximum object size.
- `object_age_days` - Number of days since creation after which objects can be deleted.
- `object_expiry_date` - Specific date from when objects can expire.
- `expired_object_delete_marker` - Cleanup object delete markers.
- `new_non_current_versions` - Number of latest non-current versions to be retained.
- `non_current_days` - Number of days after which non-current versions can be deleted.
- `after_initiation_days` - Number of days of initiation after which uploads can be aborted.

#### Related ONTAP commands

- `vserver object-store-server bucket lifecycle-management-rule create`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | Unique identifier of a bucket. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

### Example request

```
{
  "bucket_name": "bucket1",
  "expiration": {
    "object_age_days": 100,
    "object_expiry_date": "2039-09-23 00:00:00 +0000"
  },
  "name": "string",
  "object_filter": {
    "prefix": "/logs",
    "size_greater_than": 10240,
    "size_less_than": 10485760,
    "tags": [
      "project1=projA",
      "project2=projB"
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92406111 + | "Lifecycle Management rule "rule1" for bucket "testbuck1" in SVM "vs0" is invalid because of the number of tags specified. The maximum number of tags supported is four."; |
| 92406113 + | "Lifecycle Management rule for bucket in Vserver is invalid because there is a mismatch in one of the filter-field. Filters must be the same for a particular rule identifier."; |
| 92406114 + | "The Expiration action requires specifying either object_expiry_date, object_age_days or expired_object_delete_marker."; |
| 92406115 + | "The NonCurrentVersionExpiration action requires either new_non_current_versions or non_current_days."; |
| 92406116 + | "The AbortIncompleteMultipartUpload action requires after_initiation_days."; |
| 92406117 + | "The "Expiration" action cannot have both an expiry date and an age."; |
| 92406118 + | "Using Lifecycle Management rules requires an effective cluster version of 9.13.1 ."; |
| 92406120 + | "The "AbortIncompleteMultipartUpload" action cannot be specified with tags."; |
| 92406121 + | ""expired_object_delete_marker" cannot be specified with "tags"." |

| Error code | Message |
| --- | --- |
| 92406122 + | ""expired_object_delete_marker" cannot be specified with "object_age_days" or "object_expiry_date".."; |
| 92406123 + | "Expiration is supported on object store volumes only, bucket "testbuck-nas" on SVM "vs0" is not an object store volume."; |
| 92406126 + | "Lifecycle Management rule "rule1" for bucket "testbuck1" in SVM "vs0" is invalid because specified tags contain one or more invalid characters. Valid characters for a tag are 0-9, A-Z, a-z, "+", "-", "=", ".", "_", ":", "/", "@", and " ". Each "=" character present in a tag key or value must be prefixed with the "" escape character."; |
| 92406127 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" is invalid because specified tag "key" length is greater than the maximum allowed length: 128."; |
| 92406128 + | "Lifecycle Management rule "rule1" for bucket "testbuck1" in SVM "vs0" has tags with duplicate keys. Verify that each tag has a unique key and then try the operation again."; |
| 92406129 + | "Lifecycle Management rule "rule1" for bucket "testbuck1" in SVM "vs0" has a prefix that is too long. The maximum length of the prefix is 1024 characters."; |
| 92406130 + | "Lifecycle Management rule "rule1" for bucket "testbuck1" in SVM "vs0" is invalid because the minimum object size of 10485760 is larger than or equal to the maximum object size of 10240."; |
| 92406131 + | "Lifecycle Management rule "testcheck2" for bucket "testbuck1" in SVM "vs0" cannot be created because "non_current_days" must be specified along with "new_non_current_versions"."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "object_age_days" to be greater than zero."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "new_non_current_versions" to be greater than zero."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "after_initiation_days" to be greater than zero."; |
| 92406133 + | "Lifecycle Management rule for bucket in Vserver is invalid. The object_expiry_date must be later than January 1, 1970."; |

| Error code | Message |
| --- | --- |
| 92406134 + | "Cannot exceed the max limit of 1000 Lifecycle Management rules per bucket."; |
| 92406135 + | "MetroCluster is configured on cluster. Object Expiration is not supported in a MetroCluster configuration."; |
| 92406136 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" is invalid. The "object_expiry_date" must be at midnight GMT."; |
| 92406139 + | "Lifecycle Management rule for bucket in Vserver with action is a stale entry. Contact technical support for assistance."; |
| 92406141 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" with action "Expiration" cannot have "expired_object_delete_marker" disabled. To disable "expired_object_delete_marker", run the "vserver object-store-server bucket lifecycle-management-rule delete" command."; |
| 92406143 + | "As part of Bucket Lifecycle, cannot create fabriclink relationship on bucket "{bucket name}" in SVM "{SVM name}". Reason : {error}."; |
| 92406144 + | "The "AbortIncompleteMultipartUpload" action cannot be specified with object size."; |
| 92406150 + | ""expired_object_delete_marker" cannot be specified with "size_less_than"."; |
| 92406148 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" cannot have "new_non_current_versions" more than 100."; |
| 92406149 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" requires an action to be specified. Retry the operation after adding an action."; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
| --- | --- | --- |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_bucket_lifecycle_rule

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |

| Name | Type | Description |
| --- | --- | --- |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 bucket lifecycle rule

DELETE `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules/{name}`

**Introduced In:** 9.13

Deletes the S3 bucket lifecycle rule configuration.

#### Related ONTAP commands

- `vserver object-store-server bucket lifecycle-management-rule delete`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| name | string | path | True | The unique identifier of the bucket lifecycle_management rule. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92406139 + | "Lifecycle Management rule for bucket in Vserver with action is a stale entry. Contact technical support for assistance."; |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve all S3 lifecycle rules for a bucket

GET `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules/{name}`

**Introduced In:** 9.13

Retrieves all S3 Lifecycle rules associated with a bucket. Note that in order to retrieve S3 bucket rule parameters, the 'fields' option should be set to '**'.

#### Related ONTAP commands

- `vserver object-store-server bucket lifecycle-management-rule show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| name | string | path | True | The name of the lifecycle management rule to be applied on the bucket. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |

| Name | Type | Description |
| --- | --- | --- |
| bucket_name | string | Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-". |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| name | string | Bucket lifecycle management rule identifier. The length of the name can range from 0 to 256 characters. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| object_filter | object_filter | Specifies a way to filter objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "abort_incomplete_multipart_upload": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "bucket_name": "bucket1",
  "expiration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "object_age_days": 100,
    "object_expiry_date": "2039-09-23 00:00:00 +0000"
  },
  "name": "string",
  "non_current_version_expiration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "object_filter": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "prefix": "/logs",
    "size_greater_than": 10240,
    "size_less_than": 10485760,
    "tags": [
      "project1=projA",
      "project2=projB"
    ]
```

```
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the S3 bucket lifecycle rule configuration

PATCH `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules/{name}`

**Introduced In:** 9.13

Updates the S3 bucket lifecycle rule configuration.

#### Important notes

- The following fields can be modified for a bucket:
- `actions` - Lifecycle Management actions associated with the rule.
- `enabled` - Lifecycle Management rule is enabled or not..
- `object_age_days` - Number of days since creation after which objects can be deleted.
- `object_expiry_date` - Specific date from when objects can expire.
- `expired_object_delete_marker` - Cleanup object delete markers.
- `new_non_current_versions` - Number of latest non-current versions to be retained.
- `non_current_days` - Number of days after which non-current versions can be deleted.
- `after_initiation_days` - Number of days of initiation after which uploads can be aborted.

#### Related ONTAP commands

- `vserver object-store-server bucket lifecycle-management-rule modify`

**Learn more**

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/rules

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| name | string | path | True | The name of the lifecycle management rule to be applied on the bucket. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs.<br>• Introduced in: 9.14 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

### Example request

```
{
  "expiration": {
    "object_age_days": 100,
    "object_expiry_date": "2039-09-23 00:00:00 +0000"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
}
```

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[s3_bucket_lifecycle_rule] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "expiration": {
        "object_age_days": 100,
        "object_expiry_date": "2039-09-23 00:00:00 +0000"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055"
    }
  ]
}
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92406114 + | "The Expiration action requires specifying either object_expiry_date, object_age_days or expired_object_delete_marker."; |
| 92406115 + | "The NonCurrentVersionExpiration action requires either new_non_current_versions or non_current_days."; |
| 92406116 + | "The AbortIncompleteMultipartUpload action requires after_initiation_days."; |
| 92406117 + | "The "Expiration" action cannot have both an expiry date and an age."; |
| 92406118 + | "Using Lifecycle Management rules requires an effective cluster version of 9.13.1 ."; |
| 92406122 + | ""expired_object_delete_marker" cannot be specified with "object_age_days" or "object_expiry_date"."; |
| 92406131 + | "Lifecycle Management rule "testcheck2" for bucket "testbuck1" in SVM "vs0" cannot be created because "non_current_days" must be specified along with "new_non_current_versions"."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "object_age_days" to be greater than zero."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "new_non_current_versions" to be greater than zero."; |
| 92406132 + | "Lifecycle Management rule "{rule}" for bucket "{bucket}" in SVM "{SVM}" requires "after_initiation_days" to be greater than zero."; |
| 92406133 + | "Lifecycle Management rule for bucket in Vserver is invalid. The object_expiry_date must be later than January 1, 1970."; |

| Error code | Message |
| --- | --- |
| 92406135 + | "MetroCluster is configured on cluster. Object Expiration is not supported in a MetroCluster configuration."; |
| 92406136 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" is invalid. The "object_expiry_date" must be at midnight GMT."; |
| 92406139 + | "Lifecycle Management rule for bucket in Vserver with action is a stale entry. Contact technical support for assistance."; |
| 92406142 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" with action "Expiration" cannot have "expired_object_delete_marker" disabled. To disable "expired_object_delete_marker", run the "vserver object-store-server bucket lifecycle-management-rule delete" command."; |
| 92406146 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" already exists. Delete the rule and then try the operation again."; |
| 92406148 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" cannot have "new_non_current_versions" more than 100."; |
| 92406149 + | "Lifecycle Management rule "rule1" for bucket "buck1" in SVM "vs0" requires an action to be specified. Retry the operation after adding an action."; |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

abort_incomplete_multipart_upload

Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags.

| Name | Type | Description |
|---|---|---|
| after_initiation_days | integer | Number of days of initiation after which uploads can be aborted. |

expiration

Specifies a way to perform expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| expired_object_delete_marker | boolean | Cleanup object delete markers. |
| object_age_days | integer | Number of days since creation after which objects can be deleted. This cannot be used along with object_expiry_date. |
| object_expiry_date | string | Specific date from when objects can expire. This cannot be used with object_age_days. |

non_current_version_expiration

Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket.

| Name | Type | Description |
|---|---|---|
| new_non_current_versions | integer | Number of latest non-current versions to be retained. |
| non_current_days | integer | Number of days after which non-current versions can be deleted. |

object_filter

Specifies a way to filter objects within a bucket.

| Name | Type | Description |
|---|---|---|
| prefix | string | A prefix that is matched against object-names within a bucket. |
| size_greater_than | integer | Size of the object greater than specified for which the corresponding lifecycle rule is to be applied. |
| size_less_than | integer | Size of the object smaller than specified for which the corresponding lifecycle rule is to be applied. |
| tags | array[string] | An array of key-value paired tags of the form {tag} or {tag=value}. |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_bucket_lifecycle_rule

Information about the lifecycle management rule of a bucket.

| Name | Type | Description |
|---|---|---|
| abort_incomplete_multipart_upload | abort_incomplete_multipart_upload | Specifies a way to perform abort_incomplete_multipart_upload action on filtered objects within a bucket. It cannot be specified with tags. |
| enabled | boolean | Specifies whether or not the associated rule is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| expiration | expiration | Specifies a way to perform expiration action on filtered objects within a bucket. |
| non_current_version_expiration | non_current_version_expiration | Specifies a way to perform non_current_version_expiration action on filtered objects within a bucket. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | Specifies the unique identifier of the bucket. |

collection_links

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 bucket snapshots

### Manage S3 bucket snapshots

#### Overview

An S3 bucket snapshot is an image of the bucket as it exists at the time when the snapshot is created.

S3 bucket snapshots can be created using scheduled snapshot policies set on the bucket, or user requested snapshot creations.

S3 bucket snapshot APIs allow you to create, delete, and retrieve snapshots.

#### Snapshot APIs

The following APIs are used to perform operations related to snapshots.

– POST /api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

– GET /api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

– GET /api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{uuid}

– DELETE /api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{uuid}

#### Examples

### Creating an S3 bucket snapshot

The POST operation is used to create an S3 bucket snapshot with the specified name.

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots:

# The call:
curl -X POST -k "https://<mgmt-
ip>/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshot
s" -H "accept: application/json" -d '{"name" : "ss1"}'

# The response:
{
"job": {
  "uuid": "af10a4d1-60b3-11ef-82cf-005056ae1130",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/af10a4d1-60b3-11ef-82cf-005056ae1130"
    }
  }
}
}

# The Job:
{
"uuid": "af10a4d1-60b3-11ef-82cf-005056ae1130",
"description": "POST /api/protocols/s3/services/148b9bbd-58d8-11ef-b7ca-
005056ae1130/buckets/2aec8270-58e3-11ef-861e-
005056ae1130/snapshots/?name=ss1",
"state": "success",
```

```
"message": "success",
"code": 0,
"start_time": "2024-08-22T14:23:52-04:00",
"end_time": "2024-08-22T14:23:54-04:00",
"svm": {
  "name": "vs0",
  "uuid": "148b9bbd-58d8-11ef-b7ca-005056ae1130",
  "_links": {
    "self": {
      "href": "/api/svm/svms/148b9bbd-58d8-11ef-b7ca-005056ae1130"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/cluster/jobs/af10a4d1-60b3-11ef-82cf-005056ae1130"
  }
}
}
```


**Retrieving S3 bucket snapshots**

The GET operation is used to retrieve all S3 bucket snapshots for a specific bucket.

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots:

# The call:
curl -X GET -k "https://<mgmt-
ip>/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshot
s" -H "accept: application/json"

# The response:
{
"records": [
  {
    "uuid": "20837456-3c8b-405a-aa3a-5300c048f87d",
    "name": "ss1"
  },
  {
    "uuid": "c67cd056-d386-477a-8378-fcc06987bedf",
    "name": "ss2"
  }
],
"num_records": 2
}
```

#### Retrieving S3 bucket snapshots and all snapshot attributes

The GET operation is used to retrieve all S3 bucket snapshots for a specific bucket along with all the snapshot attributes.

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots:

# The call:
curl -X GET -k "https://<mgmt-
ip>/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshot
s?fields=**" -H "accept: application/json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "148b9bbd-58d8-11ef-b7ca-005056ae1130"
    },
    "bucket": {
      "uuid": "2aec8270-58e3-11ef-861e-005056ae1130"
    },
    "uuid": "20837456-3c8b-405a-aa3a-5300c048f87d",
    "name": "ss1",
    "create_time": "2024-08-22T14:23:54-04:00"
  },
  {
    "svm": {
      "uuid": "148b9bbd-58d8-11ef-b7ca-005056ae1130"
    },
    "bucket": {
      "uuid": "2aec8270-58e3-11ef-861e-005056ae1130"
    },
    "uuid": "c67cd056-d386-477a-8378-fcc06987bedf",
    "name": "ss2",
    "create_time": "2024-08-22T14:30:42-04:00"
  }
],
"num_records": 2
}
```

#### Retrieving the attributes of a specific S3 bucket snapshot

The GET operation is used to retrieve the attributes of a specific S3 bucket snapshot.

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{
uuid}:

# The call:
curl -X GET -k "https://<mgmt-
ip>/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshot
s/{uuid}" -H "accept: application/json"

# The response:
{
"svm": {
  "uuid": "148b9bbd-58d8-11ef-b7ca-005056ae1130"
},
"bucket": {
  "uuid": "2aec8270-58e3-11ef-861e-005056ae1130"
},
"uuid": "c67cd056-d386-477a-8378-fcc06987bedf",
"name": "ss2",
"create_time": "2024-08-22T14:30:42-04:00"
}
```

#### Deleting an S3 bucket snapshot

The DELETE operation is used to delete an S3 bucket snapshot.

```
# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{
uuid}:

# The call:
curl -X DELETE -k "https://<mgmt-
ip>/api/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshot
s/{uuid}" -H "accept: application/json"

# The response:
{
"job": {
  "uuid": "94f97fa9-60b7-11ef-82cf-005056ae1130",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/94f97fa9-60b7-11ef-82cf-005056ae1130"
    }
  }
}
```

```
}

# The Job:
{
"uuid": "94f97fa9-60b7-11ef-82cf-005056ae1130",
"description": "DELETE /api/protocols/s3/services/148b9bbd-58d8-11ef-b7ca-
005056ae1130/buckets/2aec8270-58e3-11ef-861e-
005056ae1130/snapshots/c67cd056-d386-477a-8378-fcc06987bedf",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2024-08-22T14:51:46-04:00",
"end_time": "2024-08-22T14:51:47-04:00",
"svm": {
  "name": "vs0",
  "uuid": "148b9bbd-58d8-11ef-b7ca-005056ae1130",
  "_links": {
    "self": {
      "href": "/api/svm/svms/148b9bbd-58d8-11ef-b7ca-005056ae1130"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/cluster/jobs/94f97fa9-60b7-11ef-82cf-005056ae1130"
  }
}
}
```

### Retrieve a collection of S3 bucket snapshots

GET `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots`

**Introduced In:** 9.16

Retrieves a collection of S3 bucket snapshots.

#### Related ONTAP commands

- `vserver object-store-server bucket snapshot show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| create_time | string | query | False | Filter by create_time |
| name | string | query | False | Filter by name |
| svm.name | string | query | False | Filter by svm.name |
| uuid | string | query | False | Filter by uuid |
| bucket_uuid | string | query | False | Filter by bucket_uuid |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_bucket_snapshot] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "bucket_uuid": "2aec8270-58e3-11ef-861e-005056ae1130",
      "create_time": "2024-08-21 20:18:04 -0400",
      "name": "snap1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "522b29a1-3b26-11e9-bd58-0050568ea321"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_bucket_snapshot

Information about an S3 bucket snapshot.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| bucket_uuid | string | The unique identifier of the bucket. |
| create_time | string | Creation time of the snapshot. It is the storage unit access time when the snapshot was created. |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the snapshot. Snapshot names must start with a lowercase letter, a number, or a hyphen, must end with a lowercase letter or a number, and cannot exceed 30 characters. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | The UUID of the snapshot in the bucket that uniquely identifies the snapshot in that bucket. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an S3 bucket snapshot**

POST `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots`

**Introduced In:** 9.16

Creates an S3 bucket snapshot.

**Required properties**

- `name` - Name of the S3 bucket snapshot to be created.

#### Related ONTAP commands

- `vserver object-store-server bucket snapshot create`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | Unique identifier of a bucket. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| bucket_uuid | string | The unique identifier of the bucket. |
| create_time | string | Creation time of the snapshot. It is the storage unit access time when the snapshot was created. |
| name | string | The name of the snapshot. Snapshot names must start with a lowercase letter, a number, or a hyphen, must end with a lowercase letter or a number, and cannot exceed 30 characters. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | The UUID of the snapshot in the bucket that uniquely identifies the snapshot in that bucket. |

### Example request

```
{
  "bucket_uuid": "2aec8270-58e3-11ef-861e-005056ae1130",
  "create_time": "2024-08-21 20:18:04 -0400",
  "name": "snap1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "522b29a1-3b26-11e9-bd58-0050568ea321"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92406249 + | "Failed to create S3 bucket snapshot "{snapshot name}" for S3 bucket "{bucket name}" in SVM "{svm name}". Reason: {Reason for failure} "; |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_bucket_snapshot

Information about an S3 bucket snapshot.

| Name | Type | Description |
|---|---|---|
| bucket_uuid | string | The unique identifier of the bucket. |
| create_time | string | Creation time of the snapshot. It is the storage unit access time when the snapshot was created. |
| name | string | The name of the snapshot. Snapshot names must start with a lowercase letter, a number, or a hyphen, must end with a lowercase letter or a number, and cannot exceed 30 characters. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | The UUID of the snapshot in the bucket that uniquely identifies the snapshot in that bucket. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 bucket snapshot

DELETE `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{uuid}`

**Introduced In:** 9.16

Deletes S3 bucket snapshot.

#### Related ONTAP commands

- `vserver object-store-server bucket snapshot delete`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| uuid | string | path | True | The unique identifier of the S3 bucket snapshot. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error code | Message |
|---|---|
| 92405861 + | "The specified SVM UUID or bucket UUID does not exist."; |
| 92406250 + | "Failed to delete S3 bucket snapshot "{snapshot name}" for S3 bucket "{bucket name}" in SVM "{svm name}". Reason: {Reason for failure} "; |
| 92406278 + | "Failed to delete S3 bucket snapshot with UUID "{snapshot uuid}" for S3 bucket "{bucket name}" in SVM "{svm name}". Reason: {Reason for failure} "; |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve details of a specific S3 bucket snapshot

GET `/protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots/{uuid}`

**Introduced In:** 9.16

Retrieves details of a specific S3 bucket snapshot.

#### Related ONTAP commands

- `vserver object-store-server bucket snapshot show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/buckets/{s3_bucket.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3_bucket.uuid | string | path | True | The unique identifier of the bucket. |
| uuid | string | path | True | The unique identifier of the S3 bucket snapshot. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| bucket_uuid | string | The unique identifier of the bucket. |
| create_time | string | Creation time of the snapshot. It is the storage unit access time when the snapshot was created. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the snapshot. Snapshot names must start with a lowercase letter, a number, or a hyphen, must end with a lowercase letter or a number, and cannot exceed 30 characters. |
| svm | svm | Specifies the name of the SVM where this bucket exists. |
| uuid | string | The UUID of the snapshot in the bucket that uniquely identifies the snapshot in that bucket. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bucket_uuid": "2aec8270-58e3-11ef-861e-005056ae1130",
  "create_time": "2024-08-21 20:18:04 -0400",
  "name": "snap1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "522b29a1-3b26-11e9-bd58-0050568ea321"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

Specifies the name of the SVM where this bucket exists.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage S3 service group configuration

## Manage S3 service group configuration

### Overview

An S3 group consists of one or many users. Policies are attached to the S3 group to have access control over S3 resources at group level.

### Examples

## Retrieving all fields for all S3 groups of an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/groups:

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-
0050568ea123/groups?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "comment": "Admin group",
    "id": 5,
    "name": "Admin-Group",
    "policies": [
      { "name": "Policy1" },
      { "name": "Policy2" },
      { "name": "Policy3" }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "users": [
      { "name": "User1" },
      { "name": "User2" },
      { "name": "User3" }
    ]
  },
  {
    "comment": "Admin group",
    "id": 6,
```

```
    "name": "Admin-Group1",
    "policies": [
      { "name": "Policy1" },
      { "name": "Policy2" },
      { "name": "Policy3" }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "users": [
      { "name": "User1" },
      { "name": "User2" },
      { "name": "User6" }
    ]
  }
],
"num_records": 2
}
```


**Retrieving the specified group in the SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/groups/5?fields=*" -H "accept: application/json"

# The response:
{
"comment": "Admin group",
"id": 5,
"name": "Admin-Group",
"policies": [
  { "name": "Policy1" },
  { "name": "Policy2" },
  { "name": "Policy3" }
],
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"users": [
  { "name": "User1" },
  { "name": "User2" },
  { "name": "User3" }
]
}
```

**Creating an S3 group for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/groups

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/groups?return_records=true" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>" -H "Content-
Type: application/json" -d "{ \"comment\": \"Admin group\", \"name\":
\"Admin-Group\", \"policies\": [ { \"name\": \"Policy1\" }, { \"name\":
\"Policy2\" }, { \"name\": \"Policy3\" } ], \"users\": [ { \"name\":
\"User1\" }, { \"name\": \"User2\" }, { \"name\": \"User3\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 289
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "comment": "Admin group",
    "id": 5,
    "name": "Admin-Group",
    "policies": [
      { "name": "Policy1" },
      { "name": "Policy2" },
      { "name": "Policy3" }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "users": [
      { "name": "User1" },
      { "name": "User2" },
      { "name": "User3" }
    ]
  }
]
}
```

#### Updating an S3 group for an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/groups/5" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"comment\": \"Admin group\", \"name\": \"Admin-Group\",
\"policies\": [ { \"name\": \"Policy1\" } ], \"users\": [ { \"name\":
\"user-1\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
```

#### Deleting an S3 group for a specified SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/groups/5" -H "accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
```

#### Retrieve an S3 group configuration

GET `/protocols/s3/services/{svm.uuid}/groups`

**Introduced In:** 9.8

Retrieves the S3 group's SVM configuration.

#### Related ONTAP commands

- `vserver object-store-server group show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| policies.name | string | query | False | Filter by policies.name<br>• maxLength: 128<br>• minLength: 1 |
| name | string | query | False | Filter by name<br>• maxLength: 128<br>• minLength: 1 |
| users.name | string | query | False | Filter by users.name<br>• maxLength: 64<br>• minLength: 1 |
| id | integer | query | False | Filter by id |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_group] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "comment": "Admin group",
      "id": 5,
      "name": "Admin-Group",
      "policies": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "Policy1"
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "user-1"
        }
      ]
```

```
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

policies

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

s3_group

This is a container for S3 user groups.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | Specifies the list of users who belong to the group. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an S3 group configuration

POST `/protocols/s3/services/{svm.uuid}/groups`

**Introduced In:** 9.8

Creates the S3 group configuration.

#### Important notes

- Each SVM can have one or more s3 group configurations.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the user configuration.
- `name` - Group name that is to be created.
- `users` - List of users to be added into the group.
- `policies` - List of policies are to be attached to this group.

#### Recommended optional properties

- `comment` - Short description about the S3 Group.

#### Related ONTAP commands

- `vserver object-store-server group create`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
| --- | --- | --- |
| users | array[users] | Specifies the list of users who belong to the group. |

#### Example request

```
{
  "comment": "Admin group",
  "id": 5,
  "name": "Admin-Group",
  "policies": [
    {
      "name": "Policy1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "user-1"
    }
  ]
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[s3_group] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "comment": "Admin group",
      "id": 5,
      "name": "Admin-Group",
      "policies": [
        {
          "name": "Policy1"
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        {
          "name": "user-1"
        }
      ]
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92406207 + | An object store group cannot have an empty list of users. |

| Error Code | Description |
| --- | --- |
| 92405930 + | Group name "{group name}" is not valid. Group names must have between 1 and 128 characters. |
| 92405931 + | Group name "{group name}" contains invalid characters. Valid characters for a group name are 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-". |
| 92405932 + | Creating an object store group requires an effective cluster version of 9.8 or later. |
| 92405933 + | Failed to create group for SVM. Reason: "{reason of failure}". |
| 92405934 + | Users {user name} specified in the user list do not exist for SVM "{vserver name}". Use the "object-store-server user create" command to create a user. |
| 92405935 + | Internal Error. Unable to generate unique ID for the group. |
| 92405936 + | Query characters such as * are not supported in the user list. |
| 92405937 + | Policies {policy name} specified in the policy list do not exist for SVM "{vserver name}". |
| 92405942 + | Group name "{group name}" already exists for SVM "{vserver name}". |
| 92405858 + | Failed to "create" the "group" because the operation is only supported on data SVMs. |
| 262222 + | User name is present more than once in the users field. |
| 92406154 + | Users {user} specified in the user list do not exist for SVM "vs1". Use the "/protocols/s3/services/{svm.uuid}/users" endpoint to create a user. |
| 92406307 + | Cannot add object store user "user1" in SVM "vs1" to the group. Reason: The user is disabled, because it does not have keys generated. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

policies

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

s3_group

This is a container for S3 user groups.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | Specifies the list of users who belong to the group. |

collection_links

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 group configuration for an SVM

DELETE `/protocols/s3/services/{svm.uuid}/groups/{id}`

**Introduced In:** 9.8

Deletes the S3 group configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server group delete`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | string | path | True | Group identifier that identifies the unique group. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an S3 group configuration for an SVM

GET `/protocols/s3/services/{svm.uuid}/groups/{id}`

**Introduced In:** 9.8

Retrieves the S3 group configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server group show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | string | path | True | Group identifier that identifies the unique group. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | Specifies the list of users who belong to the group. |

#### Example response


```
{
  "comment": "Admin group",
  "id": 5,
  "name": "Admin-Group",
  "policies": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "Policy1"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "user-1"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

policies

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an S3 group configuration for an SVM

PATCH `/protocols/s3/services/{svm.uuid}/groups/{id}`

**Introduced In:** 9.8

Updates the S3 group configuration of an SVM.

#### Important notes

- The following fields can be modified for a group:
- `name` - Group name that needs to be modified.
- `users` - List of users present in the group.
- `policies` - List of policies to be attached to this group.

#### Recommended optional properties

- `comment` - Short description about the S3 Group.

#### Related ONTAP commands

- `vserver object-store-server group modify`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/groups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| id | integer | path | True | Group identifier that identifies the unique group. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

| Name | Type | Description |
|---|---|---|
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | Specifies the list of users who belong to the group. |

**Example request**

```
{
  "comment": "Admin group",
  "id": 5,
  "name": "Admin-Group",
  "policies": [
    {
      "name": "Policy1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "user-1"
    }
  ]
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405896 + | Users list cannot be empty. |
| 92405930 + | Group name "{group name}" is not valid. Group names must have between 1 and 128 characters. |
| 92405931 + | Group name "{group name}" contains invalid characters. Valid characters for a group name are 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-". |
| 92405934 + | Users {user name} specified in the user list do not exist for SVM "{vserver name}". Use the "object-store-server user create" command to create a user. |
| 92405936 + | Query characters such as * are not supported in the user list. |
| 92405937 + | Policies {policy name} specified in the policy list do not exist for SVM "{vserver name}". |
| 262222 + | User name is present more than once in the users field. |
| 2621706 + | The specified "{svm.uuid}" and "{svm.name}" refer to different SVMs. |
| 92406154 + | Users {user} specified in the user list do not exist for SVM "vs1". Use the "/protocols/s3/services/{svm.uuid}/users" endpoint to create a user. |
| 92406307 + | Cannot add object store user "user1" in SVM "vs1" to the group. Reason: The user is disabled, because it does not have keys generated. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

policies

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

users

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

s3_group

This is a container for S3 user groups.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Can contain any additional information about the group being created or modified. |
| id | integer | Specifies a unique Group ID used to identify a particular group. This parameter should not be specified in the POST or PATCH method. A Group ID is automatically generated and it is retrieved using the GET method. Group ID is SVM-scoped. |
| name | string | Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| policies | array[policies] | Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| users | array[users] | Specifies the list of users who belong to the group. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve S3 protocol historical performance metrics

GET `/protocols/s3/services/{svm.uuid}/metrics`

**Introduced In:** 9.8

Retrieves historical performance metrics for the S3 protocol of an SVM.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| duration | string | query | False | Filter by duration |
| status | string | query | False | Filter by status |
| svm.uuid | string | path | True | Unique identifier of the SVM. |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```


### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8586226 | UUID not found. |
| 8585947 | There are no entries matching your query. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 service policies

### Manage S3 service policies

#### Overview

An S3 policy is an object that when associated with a resource, defines their permissions. Buckets and objects are defined as resources. Policies are used to manage access to these resources.

#### Examples

### Retrieving all fields for all S3 policies of an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/policies:

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-
0050568ea123/policies?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
"num_records": 2,
"records": [
  {
```

```
    "comment": "S3 policy.",
    "name": "Policy1",
    "statements": [
      {
        "actions": [
          "*"
        ],
        "effect": "allow",
        "index": 0,
        "resources": [
          "bucket1",
          "bucket1/*"
        ],
        "sid": "FullAccessToBucket1"
      },
      {
        "actions": [
          "DeleteObject"
        ],
        "effect": "deny",
        "index": 1,
        "resources": [
          "*"
        ],
        "sid": "DenyDeleteObjectAccessToAllResources"
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
  },
  {
    "comment": "S3 policy 2.",
    "name": "Policy2",
    "statements": [
      {
        "actions": [
          "GetObject"
        ],
        "effect": "allow",
        "index": 3,
        "resources": [
          "*"
        ],
        "sid": "AllowGetObjectAccessToAllResources"
```

```
      },
      {
        "actions": [
          "*"
        ],
        "effect": "deny",
        "index": 3,
        "resources": [
          "*"
        ],
        "sid": "DenyAccessToAllResources"
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
  }
]
}
```


**Retrieving the specified policy in the SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/policies/{name}:

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/policies/Policy1?fields=*" -H "accept:
application/json"

# The response:
{
"comment": "S3 policy.",
"name": "Policy1",
"statements": [
  {
    "actions": [
      "GetObject",
      "PutObject",
      "DeleteObject",
      "ListBucket",
      "ListMyBuckets",
      "ListBucketMultipartUploads",
      "ListMultipartUploadParts",
      "GetObjectTagging",
      "PutObjectTagging",
      "DeleteObjectTagging",
      "GetBucketVersioning",
      "PutBucketVersioning"
    ],
    "effect": "deny",
    "index": 0,
    "resources": [
      "*"
    ],
    "sid": "DenyAccessToAllResources"
  }
],
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
```

**Creating an S3 policy for an SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/policies?return_records=true" -H "accept:
application/json" -H "authorization: Basic <CREDENTIALS>" -H "Content-
Type: application/json" -d "{ \"comment\": \"S3 policy.\", \"name\":
\"Policy1\", \"statements\": [ { \"actions\": [ \"ListBucket\",
\"ListMyBuckets\", \"CreateBucket\", \"DeleteBucket\" ], \"effect\":
\"allow\", \"resources\": [ \"*\" ], \"sid\":
\"AllowListAccessToAllResources\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 289
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "comment": "S3 policy.",
    "name": "Policy1",
    "statements": [
      {
        "actions": [
          "ListBucket",
          "ListMyBuckets"
        ],
        "effect": "allow",
        "index": 5,
        "resources": [
          "*"
        ],
        "sid": "AllowListAccessToAllResources"
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
```

```
]
}
```

#### Updating an S3 policy for an SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/policies/{name}:

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/policies/Policy1" -H "accept: application/json" -H
"authorization: Basic <CREDENTIALS>" -H "Content-Type: application/json"
-d "{ \"comment\": \"S3 policy.\", \"statements\": [ { \"actions\": [
\"GetObject\", \"PutObject\", \"DeleteObject\", \"ListBucket\",
\"ListMyBuckets\", \"CreateBucket\", \"DeleteBucket\"], \"effect\":
\"allow\", \"resources\": [ \"bucket1\", \"bucket1/*\" ], \"sid\":
\"FullAccessToAllResources\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
```

#### Deleting an S3 policy for a specified SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/policies/{name}:

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/policies/Policy1" -H "accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
```

### Retrieve the S3 policies in an SVM configuration

GET `/protocols/s3/services/{svm.uuid}/policies`

**Introduced In:** 9.8

Retrieves the S3 policies SVM configuration.

#### Related ONTAP commands

- `vserver object-store-server policy show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statements.actions | string | query | False | Filter by statements.actions |
| statements.resources | string | query | False | Filter by statements.resources |
| statements.effect | string | query | False | Filter by statements.effect |
| statements.index | integer | query | False | Filter by statements.index |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statements.sid | string | query | False | Filter by statements.sid<br>• maxLength: 256<br>• minLength: 0 |
| name | string | query | False | Filter by name<br>• maxLength: 128<br>• minLength: 1 |
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| svm.name | string | query | False | Filter by svm.name |
| read-only | boolean | query | False | Filter by read-only |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_policy] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "comment": "S3 policy.",
      "name": "Policy1",
      "statements": [
        {
          "actions": [
            "*"
          ],
          "effect": "allow",
          "index": 0,
          "resources": [
            "bucket1",
            "bucket1/*"
          ],
          "sid": "FullAccessToBucket1"
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

collection_links

| Name | Type | Description |
| --- | --- | --- |
| next | href | |
| self | href | |

s3_policy_statement

Specifies information about a single access policy statement.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:<br>• GetObject - retrieves objects from a bucket.<br>• PutObject - puts objects in a bucket.<br>• DeleteObject - deletes objects from a bucket.<br>• ListBucket - lists the objects in a bucket.<br>• GetBucketAcl - retrieves the access control list (ACL) of a bucket.<br>• GetObjectAcl - retrieves the access control list (ACL) of an object.<br>• ListAllMyBuckets - lists all of the buckets in a server.<br>• ListBucketMultipartUploads - lists the multipart uploads in progress for a bucket.<br>• ListMultipartUploadParts - lists the parts in a multipart upload.<br>• CreateBucket - creates a new bucket.<br>• DeleteBucket - deletes an existing bucket.<br>• GetObjectTagging - retrieves the tag set of an object.<br>• PutObjecttagging - sets the tag set for an object.<br>• DeleteObjectTagging - deletes the tag set of an object.<br>• GetBucketLocation - retrieves the location of a bucket.<br>• GetBucketVersioning - retrieves the versioning configuration of a bucket.<br>• PutBucketVersioning - modifies the versioning configuration of a bucket.<br>ListBucketVersions - lists the object versions in a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| index | integer | Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated. It is not retrieved in the GET method. |
| resources | array[string] | |
| sid | string | Specifies the statement identifier which contains additional information about the statement. |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_policy

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine

whether the request is allowed or denied.

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". It cannot be specified in a PATCH method. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an S3 policy configuration

POST `/protocols/s3/services/{svm.uuid}/policies`

**Introduced In:** 9.8

Creates the S3 policy configuration.

#### Important notes

- Each SVM can have one or more s3 policy configurations.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the s3 policy configuration.
- `name` - Policy name that is to be created.

#### Recommended optional properties

- `comment` - Short description about the S3 policy.
- `statements.effect` - Indicates whether to allow or deny access.
- `statements.actions` - List of actions that can be allowed or denied access. Example: GetObject, PutObject, DeleteObject, ListBucket, ListMyBuckets, ListBucketMultipartUploads, ListMultipartUploadParts, CreateBucket, DeleteBucket, GetObjectTagging, PutObjectTagging, DeleteObjectTagging, GetBucketVersioning, PutBucketVersioning.
- `statements.resources` - Buckets or objects that can be allowed or denied access.
- `statements.sid` - Statement identifier providing additional information about the statement.

#### Related ONTAP commands

- `vserver object-store-server policy create`
- `vserver object-store-server policy add-statement`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". It cannot be specified in a PATCH method. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
    {
      "actions": [
        "*"
      ],
      "effect": "allow",
      "index": 0,
      "resources": [
        "bucket1",
        "bucket1/*"
      ],
      "sid": "FullAccessToBucket1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[s3_policy] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "comment": "S3 policy.",
      "name": "Policy1",
      "statements": [
        {
          "actions": [
            "*"
          ],
          "effect": "allow",
          "index": 0,
          "resources": [
            "bucket1",
            "bucket1/*"
          ],
          "sid": "FullAccessToBucket1"
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 92405906 + | The specified action name is invalid. |
| 92405947 + | Creating an object store server policy or statement requires an effective cluster version of 9.8 or later. |
| 92405948 + | Policy name is not valid. Policy names must have between 1 and 128 characters. |
| 92405949 + | Policy name contains invalid characters. Valid characters: 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-". |
| 92405950 + | Policy name already exists for SVM. |
| 92405954 + | Policy name is reserved for read-only policies. Cannot be used for custom policy creation. |
| 92405963 + | Failed to create policy statements for policy. Reason: "{reason of failure}". Resolve all issues and retry the operation. |
| 92405863 + | Failed to create s3 policy statements. Reason: "{reason of failure}". Valid ways to specify a resource are "__", "{bucket-name}", "{bucket-name}/…/…".". Resolve all the issues and retry the operation. |

**Definitions**

**See Definitions**

s3_policy_statement

Specifies information about a single access policy statement.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:<br>• GetObject - retrieves objects from a bucket.<br>• PutObject - puts objects in a bucket.<br>• DeleteObject - deletes objects from a bucket.<br>• ListBucket - lists the objects in a bucket.<br>• GetBucketAcl - retrieves the access control list (ACL) of a bucket.<br>• GetObjectAcl - retrieves the access control list (ACL) of an object.<br>• ListAllMyBuckets - lists all of the buckets in a server.<br>• ListBucketMultipartUploads - lists the multipart uploads in progress for a bucket.<br>• ListMultipartUploadParts - lists the parts in a multipart upload.<br>• CreateBucket - creates a new bucket.<br>• DeleteBucket - deletes an existing bucket.<br>• GetObjectTagging - retrieves the tag set of an object.<br>• PutObjecttagging - sets the tag set for an object.<br>• DeleteObjectTagging - deletes the tag set of an object.<br>• GetBucketLocation - retrieves the location of a bucket.<br>• GetBucketVersioning - retrieves the versioning configuration of a bucket.<br>• PutBucketVersioning - modifies the versioning configuration of a bucket.<br>ListBucketVersions - lists the object versions in a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| index | integer | Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated. It is not retrieved in the GET method. |
| resources | array[string] | |
| sid | string | Specifies the statement identifier which contains additional information about the statement. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_policy

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine

whether the request is allowed or denied.

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". It cannot be specified in a PATCH method. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

collection_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 policy configuration

DELETE `/protocols/s3/services/{svm.uuid}/policies/{name}`

**Introduced In:** 9.8

Deletes the S3 policy configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server policy delete`
- `vserver object-store-server policy delete-statement`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Policy name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405858 + | Failed to "delete" the "policy" because the operation is only supported on data SVMs. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an S3 policy configuration

GET `/protocols/s3/services/{svm.uuid}/policies/{name}`

**Introduced In:** 9.8

Retrieves the S3 policy configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server policy show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Policy name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| name | string | Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". It cannot be specified in a PATCH method. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
    {
      "actions": [
        "*"
      ],
      "effect": "allow",
      "index": 0,
      "resources": [
        "bucket1",
        "bucket1/*"
      ],
      "sid": "FullAccessToBucket1"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

s3_policy_statement

Specifies information about a single access policy statement.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:<br>• GetObject - retrieves objects from a bucket.<br>• PutObject - puts objects in a bucket.<br>• DeleteObject - deletes objects from a bucket.<br>• ListBucket - lists the objects in a bucket.<br>• GetBucketAcl - retrieves the access control list (ACL) of a bucket.<br>• GetObjectAcl - retrieves the access control list (ACL) of an object.<br>• ListAllMyBuckets - lists all of the buckets in a server.<br>• ListBucketMultipartUploads - lists the multipart uploads in progress for a bucket.<br>• ListMultipartUploadParts - lists the parts in a multipart upload.<br>• CreateBucket - creates a new bucket.<br>• DeleteBucket - deletes an existing bucket.<br>• GetObjectTagging - retrieves the tag set of an object.<br>• PutObjecttagging - sets the tag set for an object.<br>• DeleteObjectTagging - deletes the tag set of an object.<br>• GetBucketLocation - retrieves the location of a bucket.<br>• GetBucketVersioning - retrieves the versioning configuration of a bucket.<br>• PutBucketVersioning - modifies the versioning configuration of a bucket.<br>ListBucketVersions - lists the object versions in a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| index | integer | Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated. It is not retrieved in the GET method. |
| resources | array[string] | |
| sid | string | Specifies the statement identifier which contains additional information about the statement. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an S3 policy configuration

PATCH `/protocols/s3/services/{svm.uuid}/policies/{name}`

**Introduced In:** 9.8

Updates the S3 policy configuration of an SVM.

#### Important notes

- The following fields can be modified for a policy:
  - `comment` - Any information related to the policy.
  - `statements` - Specifies the array of policy statements.

#### Related ONTAP commands

- `vserver object-store-server policy modify`
- `vserver object-store-server policy modify-statement`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Policy name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "comment": "S3 policy.",
  "statements": [
    {
      "actions": [
        "*"
      ],
      "effect": "allow",
      "index": 0,
      "resources": [
        "bucket1",
        "bucket1/*"
      ],
      "sid": "FullAccessToBucket1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405906 + | The specified action name is invalid. |
| 92405963 + | Failed to create s3 policy statements "{policy name}". Reason: "{reason of failure}". Resolve all issues and retry the operation. |
| 92405953 + | Object store server read-only policies do not support create, modify, delete, add-statement, remove-statement and modify-statement operations. |

| Error Code | Description |
| --- | --- |
| 92406075 + | Failed to modify policy statement for policy "{policy name}". Reason: "{reason of failure}". Valid ways to specify a resource are "__", "{bucket-name}", "{bucket-name}/…/…".". |

**Definitions**

**See Definitions**

s3_policy_statement

Specifies information about a single access policy statement.

| Name | Type | Description |
| --- | --- | --- |
| actions | array[string] | For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:<br>• GetObject - retrieves objects from a bucket.<br>• PutObject - puts objects in a bucket.<br>• DeleteObject - deletes objects from a bucket.<br>• ListBucket - lists the objects in a bucket.<br>• GetBucketAcl - retrieves the access control list (ACL) of a bucket.<br>• GetObjectAcl - retrieves the access control list (ACL) of an object.<br>• ListAllMyBuckets - lists all of the buckets in a server.<br>• ListBucketMultipartUploads - lists the multipart uploads in progress for a bucket.<br>• ListMultipartUploadParts - lists the parts in a multipart upload.<br>• CreateBucket - creates a new bucket.<br>• DeleteBucket - deletes an existing bucket.<br>• GetObjectTagging - retrieves the tag set of an object.<br>• PutObjecttagging - sets the tag set for an object.<br>• DeleteObjectTagging - deletes the tag set of an object.<br>• GetBucketLocation - retrieves the location of a bucket.<br>• GetBucketVersioning - retrieves the versioning configuration of a bucket.<br>• PutBucketVersioning - modifies the versioning configuration of a bucket.<br>ListBucketVersions - lists the object versions in a bucket. |

| Name | Type | Description |
| --- | --- | --- |
| effect | string | Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access. |
| index | integer | Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated. It is not retrieved in the GET method. |
| resources | array[string] | |
| sid | string | Specifies the statement identifier which contains additional information about the statement. |

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_policy

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine

whether the request is allowed or denied.

| Name | Type | Description |
|---|---|---|
| comment | string | Can contain any additional information about the S3 policy. |
| read-only | boolean | Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method. |
| statements | array[s3_policy_statement] | Specifies the policy statements. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage S3 service users

### Manage S3 service users

#### Overview

An S3 user account is created on the S3 server. Buckets that are created for the server are associated with that user (as the owner of the buckets).
The creation of the user account involves generating a pair of keys "access" and "secret".
These keys are shared with clients (by the administrator out of band) who want to access the S3 server. The

access_key is sent in the request and it identifies the user performing the operation. The client or server never send the secret_key over the wire.
Only the access_key can be retrieved from a GET operation. The secret_key along with the access_key is returned from a POST operation and from a PATCH operation if the administrator needs to regenerate the keys.
If the user is part of active-directory, the user name takes the format "user@fully_qualified_domain_name".

#### Examples

### Retrieving S3 user configurations for a particular SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users?fields=*&return_records=true" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
        }
      }
    },
    "name": "user-1",
    "comment": "S3 user",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "key_time_to_live": "PT3H5M",
    "key_expiry_time": "2023-11-13T23:28:03+05:30",
    "id": 3,
    "keys": [
      {
        "id": 1,
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "time_to_live": "PT3H5M",
        "expiry_time": "2023-11-13T23:28:03+05:30"
      },
      {
        "id": 2,
```

```
        "access_key": "<AWS-ACCESS-KEY-ID>",
        "time_to_live": "PT6H3M2S",
        "expiry_time": "2023-11-13T23:24:41+05:30"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1"
      }
    }
  },
  {
    "svm": {
      "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
        }
      }
    },
    "name": "user-2",
    "comment": "s3-user",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "id": 2,
    "keys": [
      {
        "id": 1,
        "access_key": "<AWS-ACCESS-KEY-ID>",
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users?fields=*&return_records=true"
  }
```

```
  }
  }
```

### Retrieving the user configuration of a specific S3 user

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-1" -H "accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
    }
  }
},
"name": "user-1",
"comment": "s3-user",
"access_key": "<AWS-ACCESS-KEY-ID>",
"key_time_to_live": "P6DT1H5M",
"key_expiry_time": "2023-02-20T10:04:31Z",
"id": 3,
"keys": [
  {
    "id": 1,
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "time_to_live": "PT3H5M",
    "expiry_time": "2023-11-13T23:28:03+05:30"
  },
  {
    "id": 2,
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "time_to_live": "PT6H3M2S",
    "expiry_time": "2023-11-13T23:24:41+05:30"
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1"
  }
}
}
```

#### Creating an S3 user configuration

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"name\": \"user-1\"}"

# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 09:34:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1
Content-Length: 244
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "user-1",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>"
  }
]
}
```

#### Creating an S3 user configuration with key expiration configuration

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"comment\": \"S3 user3\",
\"key_time_to_live\": \"P6DT1H5M\", \"name\": \"user-3\"}"

# The response:
HTTP/1.1 201 Created
Date: Tue, 14 Feb 2023 08:59:31 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/protocols/s3/services/6573ac2b-ab66-11ed-b53d-
005056bb4b9b/users/user-3
Content-Length: 337
Content-Type: application/hal+json
Vary: Origin
{
"num_records": 1,
"records": [
  {
    "name": "user-3",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>",
    "key_expiry_time": "2023-06-16T12:08:38Z",
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/6573ac2b-ab66-11ed-b53d-
005056bb4b9b/users/user-3"
      }
    }
  }
]
}
```

**Creating an S3 user configuration with user keys specified**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users" -H "accept: application/json" -H "Content-
Type: application/json" -d "{  \"access_key\": \"<AWS-ACCESS-KEY-ID>\",
\"comment\": \"S3 user\",  \"key_time_to_live\": \"PT6H3M\",  \"name\":
\"user-test\",  \"secret_key\": \"<AWS-SECRET-ACCESS-KEY>\"}"

# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 09:34:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1
Content-Length: 244
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "name": "user-test",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>",
    "key_expiry_time": "2024-02-12T22:36:29+05:30"
  }
]
}
```

**Creating an S3 user configuration with a key expiration configuration and where the user is part of Active directory.**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"comment\": \"S3 user3\",
\"key_time_to_live\": \"P6DT1H5M\", \"name\": \"user-3@domain1.com\"}"

# The response:
HTTP/1.1 201 Created
Date: Tue, 14 Feb 2023 08:59:31 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/protocols/s3/services/6573ac2b-ab66-11ed-b53d-
005056bb4b9b/users/user-3%40domain1.com
Content-Length: 337
Content-Type: application/hal+json
Vary: Origin
{
"num_records": 1,
"records": [
  {
    "name": "user-3@domain1.com",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>",
    "key_expiry_time": "2023-06-16T12:08:38Z",
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/6573ac2b-ab66-11ed-b53d-
005056bb4b9b/users/user-3%40domain1.com"
      }
    }
  }
]
}
```

**Regenerating first key for a specific S3 user for the specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-2?regenerate_keys=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ }"

# The response:
HTTP/1.1 200 OK
Date: Fri, 31 May 2019 09:55:45 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 391
Content-Type: application/hal+json
{
"num_records": 1,
"records": [
  {
    "name": "user-2",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>",
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-2"
      }
    }
  }
]
}
```

**Regenerating keys and setting new expiry configuration for a specific S3 user for the specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-2?regenerate_keys=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"key_time_to_live\": \"PT6H3M\" }"

# The response:
HTTP/1.1 200 OK
Date: Fri, 31 May 2019 09:55:45 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 391
Content-Type: application/hal+json
{
"num_records": 1,
"records": [
  {
    "name": "user-2",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<>AWS-SECRET-ACCESS-KEY",
    "key_expiry_time": "2023-06-16T16:19:06Z",
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-2"
      }
    }
  }
]
}
```

**Regenerating keys by specifying keys in the input**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-2?regenerate_keys=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"access_key\": \"<AWS-ACCESS-KEY-ID>\",  \"comment\": \"S3 user\",
\"key_time_to_live\": \"PT6H3M\",  \"secret_key\": \"<AWS-SECRET-ACCESS-
KEY>\"}"

# The response:
HTTP/1.1 200 OK
Date: Fri, 31 May 2019 09:55:45 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 391
Content-Type: application/hal+json
{
"num_records": 1,
"records": [
  {
    "name": "user-2",
    "access_key": "<AWS-ACCESS-KEY-ID>",
    "secret_key": "<AWS-SECRET-ACCESS-KEY>",
    "key_expiry_time": "2023-06-16T16:19:06Z",
    "_links": {
      "self": {
        "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-2"
      }
    }
  }
]
}
```

**Deleting first key for a specific S3 user for a specified SVM**

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-2?delete_keys=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ }"

# The response:
HTTP/1.1 200 OK
Date: Wed, 08 Feb 2023 13:40:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Length: 3
Content-Type: application/hal+json
Vary: Origin
{
}
```

#### Deleting the specified S3 user configuration for a specified SVM

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298/users/user-2" -H "accept: application/json"
```

### Retrieve the SVM configuration for an S3 user

GET `/protocols/s3/services/{svm.uuid}/users`

**Introduced In:** 9.7

Retrieves the S3 user's SVM configuration.

#### Related ONTAP commands

- `vserver object-store-server user show`

**Learn more**

- DOC /protocols/s3/services/{svm.uuid}/users

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| comment | string | query | False | Filter by comment<br>• maxLength: 256<br>• minLength: 0 |
| key_time_to_live | string | query | False | Filter by key_time_to_live<br>• Introduced in: 9.14 |
| svm.name | string | query | False | Filter by svm.name |
| name | string | query | False | Filter by name<br>• maxLength: 64<br>• minLength: 1 |
| key_expiry_time | string | query | False | Filter by key_expiry_time<br>• Introduced in: 9.14 |
| access_key | string | query | False | Filter by access_key |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | collection_links | |
| num_records | integer | Number of records |
| records | array[s3_user] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "access_key": "<AWS-ACCESS-KEY-ID>",
      "comment": "S3 user",
      "key_expiry_time": "2024-01-01 00:00:00 +0000",
      "key_time_to_live": "PT6H3M",
      "name": "user-1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

collection_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |

| Name | Type | Description |
| --- | --- | --- |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an S3 user configuration

POST `/protocols/s3/services/{svm.uuid}/users`

**Introduced In:** 9.7

Creates the S3 user configuration.

#### Important notes

- Each SVM can have one or more user configurations.
- If the user is a member of Active directory, the user name takes the format "user@FQDN". For example, "user1@domain1.com".
- If user creation is successful, a user access_key and secret_key is returned as part of the response.
- If user keys have expiry configuration, then "key_expiry_time" is also returned as part of the response.
- User access and secret keys can be specified in the input. If not specified, keys are generated automatically.

#### Required properties

- `svm.uuid` - Existing SVM in which to create the user configuration.
- `name` - User name that is to be created.

#### Default property values

- `comment` - ""

#### Related ONTAP commands

- `vserver object-store-server user create`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |

| Name | Type | Description |
|---|---|---|
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example request**

```
{
  "access_key": "<AWS-ACCESS-KEY-ID>",
  "comment": "S3 user",
  "key_expiry_time": "2024-01-01 00:00:00 +0000",
  "key_time_to_live": "PT6H3M",
  "name": "user-1",
  "secret_key": "<AWS-SECRET-ACCESS-KEY>",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[s3_service_user_post_response] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "access_key": "<AWS-ACCESS-KEY-ID>",
      "key_expiry_time": "2024-01-01 00:00:00 +0000",
      "name": "user-1",
      "secret_key": "<AWS-SECRET-ACCESS-KEY>"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405787 + | User name "User#1" contains invalid characters. Valid characters for a user name are 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-". |
| 92405788 + | User name "User01234567890123456789012345678901234567890123456789012345678901234567890123456789012345" is not valid. User names must have between 1 and 64 characters. |
| 92405791 + | Failed to create access-key and secret-key. |
| 92405817 + | SVM "{non-data SVM name}" is not a data SVM. Specify a data SVM. |
| 92406083 + | The maximum supported value for user key expiry configuration is "1095" days. |
| 92406096 + | The user does not have permission to access the requested resource "{0}". |

| Error Code | Description |
| --- | --- |
| 92406097 + | Internal error. The operation configuration is not correct. |
| 92406196 + | The specified value for the "key_time_to_live" field cannot be greater than the maximum limit specified for the "max_key_time_to_live" field in the object store server. |
| 92406197 + | Object store user "user-2" must have a non-zero value for the "key_time_to_live" field because the maximum limit specified for the "max_key_time_to_live" field in the object store server is not zero. |
| 92406200 + | An object store user with the same access-key already exists. |
| 92406201 + | Missing access-key or secret-key. Either provide both of the keys or none. If not provided, keys are generated automatically. |
| 92406205 + | The object store user access key contains invalid characters. Valid characters are 0-9 and A-Z. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |

| Name | Type | Description |
| --- | --- | --- |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

collection_links

s3_service_user_post_response

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| key_expiry_time | string | Specifies the date and time after which the keys expire and are no longer valid. |
| name | string | The name of the user. |

| Name | Type | Description |
|---|---|---|
| secret_key | string | Specifies the secret key for the user. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an S3 user configuration

DELETE `/protocols/s3/services/{svm.uuid}/users/{name}`

**Introduced In:** 9.7

Deletes the S3 user configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server user delete`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | User name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the S3 user configuration for an SVM

GET `/protocols/s3/services/{svm.uuid}/users/{name}`

**Introduced In:** 9.7

Retrieves the S3 user configuration of an SVM.

#### Related ONTAP commands

- `vserver object-store-server user show`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | User name |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| name | string | Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".","@", and "-". |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "access_key": "<AWS-ACCESS-KEY-ID>",
  "comment": "S3 user",
  "key_expiry_time": "2024-01-01 00:00:00 +0000",
  "key_time_to_live": "PT6H3M",
  "name": "user-1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an S3 user configuration

PATCH `/protocols/s3/services/{svm.uuid}/users/{name}`

**Introduced In:** 9.7

Updates the S3 user configuration of an SVM.

#### Important notes

- User access_key and secret_key pair can be regenerated using the PATCH operation.
- User access_key and secret_key is returned in a PATCH operation if the "regenerate_keys" field is specified as true.
- If "regenerate_keys" is true and user keys have expiry configuration, then "key_expiry_time" is also returned as part of response.
- User access_key and secret_key pair can be deleted using the PATCH operation.

#### Recommended optional properties

- `regenerate_keys` - Specifies if secret_key and access_key need to be regenerated.
- `delete_keys` - Specifies if secret_key and access_key need to be deleted.
- `comment` - Any information related to the S3 user.

#### Related ONTAP commands

- `vserver object-store-server user show`
- `vserver object-store-server user regenerate-keys`
- `vserver object-store-server user delete-keys`

#### Learn more

- DOC /protocols/s3/services/{svm.uuid}/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | User name |
| regenerate_keys | boolean | query | False | Specifies whether or not to regenerate the user keys.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| delete_keys | boolean | query | False | Specifies whether or not to delete the user keys.<br>• Introduced in: 9.14<br>• Default value: |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |

| Name | Type | Description |
|---|---|---|
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "access_key": "<AWS-ACCESS-KEY-ID>",
  "comment": "S3 user",
  "key_expiry_time": "2024-01-01 00:00:00 +0000",
  "key_time_to_live": "PT6H3M",
  "secret_key": "<AWS-SECRET-ACCESS-KEY>",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[s3_service_user_post_response] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "access_key": "<AWS-ACCESS-KEY-ID>",
      "key_expiry_time": "2024-01-01 00:00:00 +0000",
      "name": "user-1",
      "secret_key": "<AWS-SECRET-ACCESS-KEY>"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 92405792 | Failed to regenerate access-key and secret-key for user. |
| 92406082 | Cannot perform "regenerate_keys" and "delete_keys" operations simultaneously on an S3 user. |
| 92406081 | The "regenerate_keys" operation on S3 User "user-2" in SVM "vs1" succeeded. However, modifying all of the other S3 user properties failed. Reason: resource limit exceeded. Retry the operation again without specifying the "regenerate_keys" parameter. |
| 92406080 | Cannot delete root user keys because there exists at least one S3 SnapMirror relationship that is using these keys. |
| 92406083 | The maximum supported value for user key expiry configuration is "1095" days. |
| 92406088 | The "key_time_to_live" parameter can only be used when the "regenerate_keys" operation is performed. |
| 92406096 | The user does not have permission to access the requested resource \"{0}\". |
| 92406097 | Internal error. The operation configuration is not correct. |

| Error Code | Description |
| --- | --- |
| 92406196 | The specified value for the "key_time_to_live" field cannot be greater than the maximum limit specified for the "max_key_time_to_live" field in the object store server. |

| 92406197 | Object store user "user-2" must have a non-zero value for the "key_time_to_live" field because the maximum limit specified for the "max_key_time_to_live" field in the object store server is not zero.
| 92406200 | An object store user with the same access-key already exists. |
| 92406201 | Missing access-key or secret-key. Either provide both of the keys or none. If not provided, keys are generated automatically. |
| 92406202 | The "delete_keys" operation must be performed without specifying the user keys. |
| 92406205 | The object store user access key contains invalid characters. Valid characters are 0-9 and A-Z. |
| 92406308 | Cannot disable user "user1" in SVM "vs1" because the user belongs to at least one object store group. Remove the user from all groups before disabling the user. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

s3_user

This is a container of S3 users.

| Name | Type | Description |
|---|---|---|
| access_key | string | Specifies the access key for the user. |
| comment | string | Can contain any additional information about the user being created or modified. |
| key_expiry_time | string | Specifies the date and time after which keys expire and are no longer valid. |

| Name | Type | Description |
| --- | --- | --- |
| key_time_to_live | string | Indicates the time period from when this parameter is specified:<br>• when creating or modifying a user or<br>• when the user keys were last regenerated, after which the user keys expire and are no longer valid.<br>• Valid format is: 'PnDTnHnMnS\|PnW'. For example, P2DT6H3M10S specifies a time period of 2 days, 6 hours, 3 minutes, and 10 seconds.<br>• If the value specified is '0' seconds, then the keys won't expire. |
| secret_key | string | Specifies the secret key for the user. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

collection_links

s3_service_user_post_response

| Name | Type | Description |
| --- | --- | --- |
| access_key | string | Specifies the access key for the user. |
| key_expiry_time | string | Specifies the date and time after which the keys expire and are no longer valid. |
| name | string | The name of the user. |
| secret_key | string | Specifies the secret key for the user. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# SAN

## ONTAP REST API SAN endpoints

### Overview

The storage area network (SAN) endpoints and objects enable you to configure, provision, and manage SAN-related objects.

---

### Fibre Channel

#### Logins

Fibre Channel logins represent connections, formed by Fibre Channel initiators, that have successfully logged in to ONTAP. This represents the Fibre Channel login on which higher-level protocols, such as Fibre Channel Protocol (FCP) and Non-Volatile Memory Express over Fibre Channel (NVMe over FC), rely.

The Fibre Channel logins REST API provides information about active Fibre Channel logins.

#### WWPN Aliases

A WWPN (world wide port name) is a unique 64-bit identifier for a Fibre Channel initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify Fibre Channel initiators using an alias, especially in larger SANs.

The WWPN alias REST API allows you to create, delete and discover aliases for WWPNs.

#### Services

A Fibre Channel Protocol (FCP) service defines the properties of the Fibre Channel Protocol target for an SVM. There can be at most one FCP service for a given SVM. An SVM's FCP service must be created before FCP initiators can login to the SVM.

The Fibre Channel Protocol (FCP) service REST API allows you to create, update, delete, and discover Fibre Channel Services for SVMs.
Fibre Channel interfaces are the logical endpoints for Fibre Channel network connections to an SVM.

---

### iSCSI

#### Credentials

An iSCSI credentials object defines the authentication credentials to be used between an iSCSI initiator and ONTAP. It identifies an authentication type, user names, and the passwords that must be used to authenticate a specific initiator.

The iSCSI credentials REST API allows you to create, update, delete, and discover iSCSI credential objects.

#### Services

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can login to the SVM.

The iSCSI service REST API allows you to create, update, delete, and discover iSCSI services for SVMs.

#### Sessions

An iSCSI session consists of one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

The iSCSI sessions REST API provides information about iSCSI initiators that have successfully logged in to ONTAP.

#### Learn More

- *IP Interfaces* found in the *networking* section. IP interfaces are the logical endpoints for iSCSI network connections to an SVM.

---

### Initiator Groups

An initiator group (igroup) is a collection of Fibre Channel WWPNs (world wide port names), iSCSI IQNs (qualified names), iSCSI EUIs (extended unique identifiers), or any combination of these, that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, a network administrator creates an initiator group containing the hosts' initiator names, and then creates a LUN map that associates the initiator group with the LUN.

The initiator group REST API allows you to create, update, delete, and discover initiator groups. It also enables you to add and remove initiators that can access the target and associated LUNs.

### Portsets

A portset is a collection of Fibre Channel Protocol and/or iSCSI network interfaces from the portset's SVM.

Portsets are used to limit the network interfaces through which an initiator can connect to mapped LUNs. When a portset is bound to an initiator group (igroup), and the initiator group is mapped to a LUN, the initiators of the initiator group can only reach the LUN through the network interfaces in the portset.

Portsets are not recommended for new configurations. With modern releases of ONTAP, it is recommended to use multiple SVMs and initiator groups with no bound portset to load balance applications over multiple ports on a node. Selective LUN mapping will automatically limit the number of visible paths to a LUN from the client host to those required for efficient access and high availability. The REST portset API is primarily intended for legacy use.

The portset REST API allows you to create, delete, and discover portsets, and to add and remove network interfaces from portsets.

A portset can be bound to one or more initiator groups. An initiator group (igroup) can be bound to at most one portset.

### LUN Maps

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the group's initiators are granted access to the LUN. The relationship between an initiator group and a LUN is many initiator groups to many LUNs.

The LUN map REST API allows you to create, delete, and discover LUN maps and manage the reporting nodes of a LUN map.

### LUNs

A LUN is the logical representation of storage in a storage area network (SAN).

The LUN REST API allows you to create, update, delete, and discover LUNs.

### VMware Virtual Volume (vVol) Bindings

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

Class `protocol_endpoint` and `vvol` LUNs support many-to-many vVol bindings. A LUN of one class can be bound to zero or more LUNs of the opposite class.

The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed.

The vVol binding REST API allows you to create, delete, and discover vVol bindings.

# View FC login information

## Retrieve FC port information

### Overview

Fibre Channel (FC) ports are the physical ports of FC adapters on ONTAP cluster nodes that can be connected to FC networks to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

The Fibre Channel port REST API allows you to discover FC ports, obtain status information for FC ports, and configure FC port properties. POST and DELETE requests are not supported. You must physically add and remove FC adapters to ONTAP nodes to create and remove ports from the ONTAP cluster.

### Performance monitoring

Performance of an FC port can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the performance of an FC port in terms of IOPS, latency, and throughput. The `metric.*` properties denote an average, whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

### Examples

## Retrieving all FC ports

```
# The API:
GET /api/network/fc/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports" -H "Accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
    "name": "0a",
    "_links": {
```

```
      "self": {
        "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
    "name": "0b",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
    "name": "0c",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
```

```
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2748-b047-11e8-9af3-005056bb838e",
    "name": "0d",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2748-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b28c2-b047-11e8-9af3-005056bb838e",
    "name": "0e",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b28c2-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
```

```
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2a7b-b047-11e8-9af3-005056bb838e",
    "name": "0f",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2a7b-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b2e2b-b047-11e8-9af3-005056bb838e",
    "name": "1b",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b2e2b-b047-11e8-9af3-
005056bb838e"
      }
    }
  }
],
"num_records": 8,
"_links": {
  "self": {
    "href": "/api/network/fc/ports"
  }
}
}
```

#### Retrieving all FC ports with state *online*

The `state` query parameter is used to perform the query.

```
# The API:
GET /api/network/fc/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports?state=online" -H
"Accept: application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
    "name": "0a",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
```

```
    "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
    "name": "0b",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-
005056bb838e"
      }
    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
        }
      }
    },
    "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
    "name": "0c",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-
005056bb838e"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/network/fc/ports?state=online"
  }
}
}
```


### Retrieving an FC port

```
# The API:
GET /api/network/fc/ports/{uuid}
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-
9af3-005056bb838e" -H "Accept: application/hal+json"

# The response:
{
"node": {
  "name": "node1",
  "uuid": "5a534a72-b047-11e8-9af3-005056bb838e",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/5a534a72-b047-11e8-9af3-005056bb838e"
    }
  }
},
"uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
"name": "0a",
"description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
Adapter, rev. 1.0.0, 8G)",
"enabled": true,
"fabric": {
  "connected": true,
  "connected_speed": 8,
  "name": "55:0e:b1:a0:20:40:80:00",
  "port_address": "52100",
  "switch_port": "ssan-g620-03:1"
},
"physical_protocol": "fibre_channel",
"speed": {
  "maximum": "8",
  "configured": "auto"
},
"state": "online",
"supported_protocols": [
  "fcp"
],
"transceiver": {
  "form_factor": "SFP",
  "manufacturer": "ACME",
  "capabilities": [
    4,
    8
  ],
  "part_number": "1000"
},
```

```
"wwnn": "50:0a:09:80:bb:83:8e:00",
"wwpn": "50:0a:09:82:bb:83:8e:00",
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
```

```
  "self": {
    "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-005056bb838e"
  }
}
}
```

### Disabling an FC port

If an active FC interface exists on an FC port, the port cannot be disabled.

```
# The API:
PATCH /api/network/fc/ports/{uuid}

# The call:
curl -X PATCH "http://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-
9af3-005056bb838e" -H "Accept: application/hal+json" -d '{ "enabled":
false }'
```

### Retrieve FC logins

GET `/network/fc/logins`

**Introduced In:** 9.6

Retrieves FC logins.

#### Related ONTAP commands

- `vserver fcp initiator show`

#### Learn more

- SAN: DOC /network/fc/logins
- NVMe: DOC /network/fc/logins

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.name | string | query | False | Filter by interface.name |
| interface.uuid | string | query | False | Filter by interface.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.wwpn | string | query | False | Filter by interface.wwpn |
| initiator.port_address | string | query | False | Filter by initiator.port_address |
| initiator.wwpn | string | query | False | Filter by initiator.wwpn |
| initiator.wwnn | string | query | False | Filter by initiator.wwnn |
| initiator.comment | string | query | False | Filter by initiator.comment<br>• Introduced in: 9.9 |
| initiator.aliases | string | query | False | Filter by initiator.aliases |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| protocol | string | query | False | Filter by protocol |
| igroups.name | string | query | False | Filter by igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| igroups.uuid | string | query | False | Filter by igroups.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fc_login] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroups": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "initiator": {
        "aliases": [
          "alias1"
        ],
        "comment": "This is an FC initiator for host 5",
        "port_address": "5060A",
        "wwnn": "2f:a0:00:a0:98:0b:56:13",
        "wwpn": "2f:a0:00:a0:98:0b:56:13"
      },
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
            }
          },
          "name": "fc_lif1",
          "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
          "wwpn": "20:00:00:50:56:b4:13:a8"
        },
        "protocol": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroups

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiator

Information about the logged in FC initiator.

| Name | Type | Description |
|---|---|---|
| aliases | array[string] | The logged in initiator world wide port name (WWPN) aliases. |
| comment | string | A comment available for use by the administrator. This is modifiable from the initiator REST endpoint directly. See PATCH /protocols/san/igroups/{igroup.uuid}/initiators/{name} . |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The port address of the initiator's FC port.<br><br>Each port in an FC switched fabric has its own unique port address for routing purposes. The port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the initiator port.<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| wwnn | string | The logged in initiator world wide node name (WWNN). |
| wwpn | string | The logged in initiator WWPN. |

interface

An FC interface.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

fc_login

A Fibre Channel (FC) login represents a connection formed by an FC initiator that has successfully logged in to ONTAP. This represents the FC login on which higher-level protocols such as Fibre Channel Protocol and NVMe over Fibre Channel (NVMe/FC) rely.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| igroups | array[igroups] | The initiator groups in which the initiator is a member. |
| initiator | initiator | Information about the logged in FC initiator. |
| interface | interface | An FC interface. |
| protocol | string | The data protocol used to perform the login. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC login

GET `/network/fc/logins/{interface.uuid}/{initiator.wwpn}`

**Introduced In:** 9.6

Retrieves an FC login.

#### Related ONTAP commands

- `vserver fcp initiator show`

#### Learn more

- SAN: DOC /network/fc/logins
- NVMe: DOC /network/fc/logins

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| interface.uuid | string | path | True | The unique identifier of the FC interface through which the initiator logged in. |
| initiator.wwpn | string | path | True | The world wide port name (WWPN) of the initiator. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| igroups | array[igroups] | The initiator groups in which the initiator is a member. |
| initiator | initiator | Information about the logged in FC initiator. |
| interface | interface | An FC interface. |
| protocol | string | The data protocol used to perform the login. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiator": {
    "aliases": [
      "alias1"
    ],
    "comment": "This is an FC initiator for host 5",
    "port_address": "5060A",
    "wwnn": "2f:a0:00:a0:98:0b:56:13",
    "wwpn": "2f:a0:00:a0:98:0b:56:13"
  },
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "fc_lif1",
    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "protocol": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The Fibre Channel login specified does not exist. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5374881 | The Fibre Channel interface specified does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroups

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiator

Information about the logged in FC initiator.

| Name | Type | Description |
|---|---|---|
| aliases | array[string] | The logged in initiator world wide port name (WWPN) aliases. |
| comment | string | A comment available for use by the administrator. This is modifiable from the initiator REST endpoint directly. See PATCH /protocols/san/igroups/{igroup.uuid}/initiators/{name} . |

| Name | Type | Description |
| --- | --- | --- |
| port_address | string | The port address of the initiator's FC port.<br><br>Each port in an FC switched fabric has its own unique port address for routing purposes. The port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the initiator port.<br><br>This is useful for obtaining statistics and diagnostic information from FC switches.<br><br>This is a hexadecimal encoded numeric value. |
| wwnn | string | The logged in initiator world wide node name (WWNN). |
| wwpn | string | The logged in initiator WWPN. |

interface

An FC interface.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage FC WWPN aliases

### Manage FC WWPN aliases

#### Overview

A worldwide port name (WWPN) is a unique 64-bit identifier for a Fibre Channel (FC) initiator. It is displayed as a 16-character hexadecimal value. SAN administrators might find it easier to identify FC initiators using an alias, especially in larger SANs.

The WWPN alias REST API allows you to create, delete, and discover aliases for WWPNs.

Multiple aliases can be created for a WWPN, but you cannot use the same alias for multiple WWPNs.

An alias can consist of up to 32 characters. Valid characters are:

- A through Z

- a through z
- numbers 0 through 9
- hyphen ("-")
- underscore ("_")
- left and right braces ("{", "}")
- period (".")

#### Examples

### Creating a WWPN alias

```
# The API:
POST /api/network/fc/wwpn-aliases

# The call:
curl -X POST "https://<mgmt-ip>/api/network/fc/wwpn-aliases" -H "Accept:
application/json" -d '{ "svm": { "name": "svm1" }, "wwpn":
"50:0a:09:82:b4:30:25:05", "alias": "alias3" }'
```

### Retrieving all properties of all WWPN aliases

The `fields` query parameter is used to request that all properties be returned.

```
# The API:
GET /api/network/fc/wwpn-aliases

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases?fields=*" -H
"Accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
        }
      }
    },
    "alias": "alias1",
```

```
    "wwpn": "20:00:00:50:56:b4:30:25",
    "_links": {
      "self": {
        "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-
005056b43025/alias1"
      }
    }
  },
  {
    "svm": {
      "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
        }
      }
    },
    "alias": "alias2",
    "wwpn": "50:0a:09:82:b4:30:25:00",
    "_links": {
      "self": {
        "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-
005056b43025/alias2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/fc/wwpn-aliases"
  }
}
}
```


---

#### Retrieving all WWPN aliases named "alias1"

The `alias` query parameter is used to specify a query for the value "alias1".

```
# The API:
GET /api/network/fc/wwpn-aliases

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases?alias=alias1"
-H "Accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
        }
      }
    },
    "alias": "alias1",
    "wwpn": "20:00:00:50:56:b4:30:25",
    "_links": {
      "self": {
        "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-
005056b43025/alias1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/fc/wwpn-aliases?alias=alias1"
  }
}
}
```

#### Retrieving a specific WWPN alias

The alias to be returned is identified by the UUID of its SVM and the alias name.

```
# The API:
GET /api/network/fc/wwpn-aliases/{svm.uuid}/{alias}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases/68589d3d-7efa-
11e8-9eed-005056b43025/alias2" -H "Accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
        }
      }
    },
    "alias": "alias2",
    "wwpn": "50:0a:09:82:b4:30:25:00",
    "_links": {
      "self": {
        "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-
005056b43025/alias1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/fc/wwpn-aliases?alias=alias1"
  }
}
}
```

### Deleting a WWPN alias

The alias to delete is identified by the UUID of its SVM and the alias name.

```
# The API:
DELETE /api/network/fc/wwpn-aliases/{svm.uuid}/{alias}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/fc/wwpn-aliases/68589d3d-
7efa-11e8-9eed-005056b43025/alias2" -H "Accept: application/hal+json"
```

### Retrieve FC WWPN aliases

GET `/network/fc/wwpn-aliases`

**Introduced In:** 9.6

Retrieves FC WWPN aliases.

#### Related ONTAP commands

- `vserver fcp wwpn-alias show`

#### Learn more

- DOC /network/fc/wwpn-aliases

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| alias | string | query | False | Filter by alias |
| wwpn | string | query | False | Filter by wwpn |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[wwpn_alias] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "alias": "host1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "wwpn": "2f:a0:00:a0:98:0b:56:13"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

wwpn_alias

A Fibre Channel (FC) world wide port name (WWPN) alias. A WWPN is a unique 64-bit identifier for an FC initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify FC initiators using an alias, especially in larger SANs.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| alias | string | The FC WWPN alias. Required in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

| Name | Type | Description |
|---|---|---|
| wwpn | string | The FC initiator WWPN. Required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an FC WWPN alias

POST `/network/fc/wwpn-aliases`

**Introduced In:** 9.6

Creates an FC WWPN alias.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the FC alias.
- `alias` - Name of the FC alias.
- `wwpn` - FC WWPN for which to create the alias.

#### Related ONTAP commands

- `vserver fcp wwpn-alias set`

#### Learn more

- DOC /network/fc/wwpn-aliases

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| alias | string | The FC WWPN alias. Required in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| wwpn | string | The FC initiator WWPN. Required in POST. |

#### Example request

```
{
  "alias": "host1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "wwpn": "2f:a0:00:a0:98:0b:56:13"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[wwpn_alias] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "alias": "host1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "wwpn": "2f:a0:00:a0:98:0b:56:13"
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The specified SVM does not exist or cannot be accessed by this user. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5374869 | The alias already exists. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

wwpn_alias

A Fibre Channel (FC) world wide port name (WWPN) alias. A WWPN is a unique 64-bit identifier for an FC initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify FC initiators using an alias, especially in larger SANs.

| Name | Type | Description |
|---|---|---|
| alias | string | The FC WWPN alias. Required in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| wwpn | string | The FC initiator WWPN. Required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an FC WWPN alias

DELETE `/network/fc/wwpn-aliases/{svm.uuid}/{alias}`

**Introduced In:** 9.6

Deletes an FC WWPN alias.

#### Related ONTAP commands

- `vserver fcp wwpn-alias remove`

#### Learn more

- DOC /network/fc/wwpn-aliases

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM. |
| alias | string | path | True | The name of FC WWPN alias. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2621462 | The specified SVM does not exist or cannot be accessed by this user. |
| 5374046 | The alias could not be found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC WWPN alias

GET `/network/fc/wwpn-aliases/{svm.uuid}/{alias}`

**Introduced In:** 9.6

Retrieves an FC WWPN alias.

#### Related ONTAP commands

- `vserver fcp wwpn-alias show`

#### Learn more

- DOC /network/fc/wwpn-aliases

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM in which the alias is found. |
| alias | string | path | True | The name of FC WWPN alias. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| alias | string | The FC WWPN alias. Required in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| wwpn | string | The FC initiator WWPN. Required in POST. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "alias": "host1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "wwpn": "2f:a0:00:a0:98:0b:56:13"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The specified SVM does not exist or cannot be accessed by this user. |
| 5374046 | The alias could not be found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage FC services for SVMs

## Manage FC services for SVMs

### Overview

A Fibre Channel Protocol (FC Protocol) service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM FC Protocol service must be created before FC Protocol initiators can log in to the SVM.

The FC Protocol service REST API allows you to create, update, delete, and discover FC services for SVMs.

### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency, and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

### Examples

## Creating an FC Protocol service for an SVM

The simplest way to create an FC Protocol service is to specify only the SVM, either by name or UUID. By default, the new FC Protocol service is enabled.

In this example, the `return_records` query parameter is used to retrieve the new FC Protocol service object in the REST response.

```
# The API:
POST /api/protocols/san/fcp/services

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/fcp/services?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
        }
      }
    },
    "enabled": true,
    "target": {
      "name": "20:00:00:50:56:bb:b2:4b"
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-
005056bbb24b"
      }
    }
  }
]
}
```

#### Retrieving FC Protocol services for all SVMs in the cluster

```
# The API:
GET /api/protocols/san/fcp/services

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services' -H 'Accept:
```

```
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-
005056bbb24b"
      }
    }
  },
  {
    "svm": {
      "uuid": "6011f874-c01a-11e8-88ed-005056bbb24b",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6011f874-c01a-11e8-88ed-005056bbb24b"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/fcp/services/6011f874-c01a-11e8-88ed-
005056bbb24b"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/fcp/services"
  }
}
```

```
}
```

### Retrieving details for a specific FC Protocol service

The FC Protocol service is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/san/fcp/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
    }
  }
},
"enabled": true,
"target": {
  "name": "20:00:00:50:56:bb:b2:4b"
},
"_links": {
  "self": {
    "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-
005056bbb24b"
  }
}
}
```

### Disabling an FC Protocol service

Disabling an FC Protocol service shuts down all active FC Protocol logins for the SVM and prevents new FC Protocol logins.

The FC Protocol service to update is identified by the UUID of its SVM.

```
# The API:
PATCH /api/protocols/san/fcp/services/{svm.uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b' -H 'Accept: application/hal+json' -d '{
"enabled": "false" }'
```

You can retrieve the FC Protocol service to confirm the change.

In this example, the `fields` query parameter is used to limit the response to the `enabled` property and FC Protocol service identifiers.

```
# The API:
GET /api/protocols/san/fcp/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b?fields=enabled' -H 'Accept:
application/hal+json'

# The response:
{
"svm": {
  "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
    }
  }
},
"enabled": false,
"_links": {
  "self": {
    "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-
005056bbb24b"
  }
}
}
```

### Deleting an FC Protocol service

The FC Protocol service must be disabled before it can be deleted.

The FC Protocol service to delete is identified by the UUID of its SVM.

```
# The API:
DELETE /api/protocols/san/fcp/services/{svm.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b' -H 'Accept: application/hal+json'
```

### Retrieve FC protocol services

GET `/protocols/san/fcp/services`

**Introduced In:** 9.6

Retrieves FC Protocol services.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver fcp show`

#### Learn more

- DOC /protocols/san/fcp/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.7 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.7 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.7 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.7 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.7 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.7 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.7 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.7 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br>• Introduced in: 9.7 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.7 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.7 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| enabled | boolean | query | False | Filter by enabled |
| target.name | string | query | False | Filter by target.name<br>• maxLength: 128<br>• minLength: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[fcp_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "statistics": {
        "iops_raw": {
```

```
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "name": "20:00:00:50:56:bb:b2:4b"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
| --- | --- | --- |
| name | string | The target name of the FC Protocol service. This is generated for the SVM during POST.<br><br>The FC Protocol target name is a world wide node name (WWNN).<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: 20:00:00:50:56:bb:b2:4b<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

fcp_service

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create an FC protocol service

POST `/protocols/san/fcp/services`

**Introduced In:** 9.6

Creates an FC Protocol service.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the FC Protocol service.

#### Related ONTAP commands

- `vserver fcp create`

#### Learn more

- DOC /protocols/san/fcp/services

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

#### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | The number of records in the response. |
| records | array[fcp_service] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "name": "20:00:00:50:56:bb:b2:4b"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115127 | The cluster lacks a valid FC Protocol license. |
| 2621462 | The supplied SVM does not exist. |
| 2621507 | The Fibre Channel Protocol is not allowed for the specified SVM. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374082 | The Fibre Channel Protocol service already exists for the SVM. |

| Error Code | Description |
| --- | --- |
| 5374092 | The Fibre Channel Protocol is not supported on the cluster hardware configuration; there are no Fibre Channel adapters. |
| 5374893 | The SVM is stopped. The SVM must be running to create a Fibre Channel Protocol service. |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| name | string | The target name of the FC Protocol service. This is generated for the SVM during POST.<br><br>The FC Protocol target name is a world wide node name (WWNN).<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: 20:00:00:50:56:bb:b2:4b<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

fcp_service

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an FC protocol service**

DELETE `/protocols/san/fcp/services/{svm.uuid}`

**Introduced In:** 9.6

Deletes an FC Protocol service. An FC Protocol service must be disabled before it can be deleted.

**Related ONTAP commands**

- `vserver fcp delete`

#### Learn more

- DOC /protocols/san/fcp/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which to delete the FC Protocol service. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 5373960 | The Fibre Channel Protocol service cannot be removed while it is enabled. |
| 5374083 | There is no Fibre Channel Protocol service for the specified SVM. |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an FC protocol service

GET `/protocols/san/fcp/services/{svm.uuid}`

**Introduced In:** 9.6

Retrieves an FC Protocol service.

#### Related ONTAP commands

- `vserver fcp show`

#### Learn more

- DOC /protocols/san/fcp/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which to retrieve the FC Protocol service. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
```

```
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 5374083 | There is no Fibre Channel Protocol service for the specified SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| name | string | The target name of the FC Protocol service. This is generated for the SVM during POST.<br><br>The FC Protocol target name is a world wide node name (WWNN).<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: 20:00:00:50:56:bb:b2:4b<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an FC protocol service

PATCH `/protocols/san/fcp/services/{svm.uuid}`

**Introduced In:** 9.6

Updates an FC Protocol service.

#### Related ONTAP commands

- `vserver fcp modify`
- `vserver fcp start`
- `vserver fcp stop`

#### Learn more

- DOC /protocols/san/fcp/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM whose FC Protocol service is to be updated. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| target | target | |

#### Example request

```
{
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2621462 | An SVM with the specified UUID does not exist. |
| 5374083 | There is no Fibre Channel Protocol service for the specified SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| name | string | The target name of the FC Protocol service. This is generated for the SVM during POST.<br><br>The FC Protocol target name is a world wide node name (WWNN).<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: 20:00:00:50:56:bb:b2:4b<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

fcp_service

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.<br><br>This is optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| target | target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for the FC protocol service of an SVM**

GET `/protocols/san/fcp/services/{svm.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for the FC Protocol service of an SVM.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| latency.other | integer | query | False | Filter by latency.other |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| status | string | query | False | Filter by status |
| duration | string | query | False | Filter by duration |
| svm.uuid | string | path | True | The unique identifier of the SVM. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "svm": {
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve historical performance metrics for the FC protocol service of an SVM for a specific time

GET `/protocols/san/fcp/services/{svm.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical performance metrics for the FC Protocol service of an SVM for a specific time.

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM. |
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "svm": {
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
}
```


### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
|---|---|---|
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

svm

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the SVM. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage SAN igroups

### Manage SAN igroups

#### Overview

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPNs), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host endpoints.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the host initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator group may contain either initiators or other initiator groups, but not both simultaneously. When a parent initiator group is mapped, it inherits all of the initiators of any initiator groups nested below it. If any nested initiator group is modified to contain different initiators, the parent initiator groups inherit the change. A parent can have many nested initiator groups and an initiator group can be nested under multiple parents. Initiators can only be added or removed from the initiator group that directly contains them. The maximum supported depth of nesting is three layers.

Best practice when using nested initiator groups is to match host hierarchies. A single initiator group should correspond to a single host. If a LUN needs to be mapped to multiple hosts, the initiator groups representing those hosts should be aggregated into a parent initiator group and the LUN should be mapped to that initiator group. For multi-ported hosts, initiators have a comment property where the port corresponding to the initiator can be documented.

The initiator group REST API allows you to create, update, delete, and discover initiator groups, and to add and remove initiators that can access the target and associated LUNs.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once. With the introduction of nestable initiator groups, best practice is to use the hierarchy such that an initiator is only a direct member of a single initiator group, and that initiator group can then be referenced by other initiator groups. This avoid needing to update multiple initiator groups when initiators change.

All initiators or nested initiator groups in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (*fcp*), iSCSI (*iscsi*), or both (*mixed*). Initiator groups within a nested hierarchy may not have conflicting protocols.

Zero or more initiators or nested initiator groups can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. Initiator groups containing other initiator groups report the aggregated list of initiators from all nested initiator groups, but modifications of the initiator list must be performed on the initiator group that directly contains the initiators. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details.

An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters.

#### Examples

### Creating an initiator group with no initiators

The example initiator group used here is for Linux iSCSI initiators only. Note that the `return_records` query parameter is used to obtain the newly created initiator group in the response.

```
# The API:
POST /api/protocols/san/igroups

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/igroups?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "igroup1",
"os_type": "linux", "protocol": "iscsi" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup1",
    "protocol": "iscsi",
    "os_type": "linux",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  }
]
}
```

#### Creating an initiator group with initiators

The example initiator group used here is for Windows. FC Protocol and iSCSI initiators are allowed. Note that the `return_records` query parameter is used to obtain the newly created initiator group in the response.

```
# The API:
```

```
POST /api/protocols/san/igroups

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/igroups?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "igroup2",
"os_type": "windows", "protocol": "mixed", "initiators": [ { "name":
"20:01:00:50:56:bb:70:72" }, { "name": "iqn.1991-05.com.ms:host1" } ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup2",
    "protocol": "mixed",
    "os_type": "windows",
    "initiators": [
      {
        "name": "20:01:00:50:56:bb:70:72",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/20:01:00:50:56:bb:70:72"
          }
        }
      },
      {
        "name": "iqn.1991-05.com.ms:host1",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/iqn.1991-05.com.ms:host1"
          }
        }
      }
```

```
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  }
]
}
```

#### Creating an initiator group with nested initiator groups

The example initiator group used here is for Windows. FC Protocol and iSCSI initiators are allowed. Note that the `return_records` query parameter is used to obtain the newly created initiator group in the response. The new initiator group is create so as to contain the initiator group created in the previous example. The initiators list reports all initiators nested below this initiator group, and note that the href link for the initiators refers to the initiator group that directly owns the initiator, not this initiator group.

```
# The API:
POST /api/protocols/san/igroups

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/igroups?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "igroup3",
"os_type": "windows", "protocol": "mixed", "igroups": [ { "name":
"igroup2" } ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7073",
    "name": "igroup3",
```

```
    "protocol": "mixed",
    "os_type": "windows",
    "igroups": [
      {
        "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
        "name": "igroup2",
        "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
      }
    ],
    "initiators": [
      {
        "name": "20:01:00:50:56:bb:70:72",
        "igroup": {
          "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
          "name": "igroup2",
          "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/20:01:00:50:56:bb:70:72"
          }
        }
      },
      {
        "name": "iqn.1991-05.com.ms:host1",
        "igroup": {
          "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
          "name": "igroup2",
          "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/iqn.1991-05.com.ms:host1"
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7073"
```

```
      }
    }
  }
]
}
```

### Retrieving all initiator groups

```
# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup1",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  },
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
```

```
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup2",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  },
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7073",
    "name": "igroup3",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7073"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups"
  }
}
}
```


#### Retrieving all properties of all initiator groups

The `fields` query parameter is used to request all initiator group properties. Note that the nested and parent initiator groups are considered expensive properties and will only be returned if explicitly requested.

```
# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/san/igroups?fields=*,igroups,parent_igroups' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup1",
    "protocol": "iscsi",
    "os_type": "linux",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  },
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup2",
    "parent_igroups": [
```

```
      {
        "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7073",
        "name": "igroup3",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7073"
          }
        }
      }
    ],
    "protocol": "mixed",
    "os_type": "windows",
    "initiators": [
      {
        "name": "20:01:00:50:56:bb:70:72",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/20:01:00:50:56:bb:70:72"
          }
        }
      },
      {
        "name": "iqn.1991-05.com.ms:host1",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/iqn.1991-05.com.ms:host1"
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  },
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
```

```
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7073",
    "name": "igroup3",
    "protocol": "mixed",
    "os_type": "windows",
    "igroups": [
      {
        "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
        "name": "igroup2",
        "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
      }
    ],
    "initiators": [
      {
        "name": "20:01:00:50:56:bb:70:72",
        "igroup": {
          "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
          "name": "igroup2",
          "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/20:01:00:50:56:bb:70:72"
          }
        }
      },
      {
        "name": "iqn.1991-05.com.ms:host1",
        "igroup": {
          "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
          "name": "igroup2",
          "_links": { "self": { "href":
"/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072" } }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7072/initiators/iqn.1991-05.com.ms:host1"
          }
        }
```

```
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-
005056bb7073"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups?fields=*"
  }
}
}
```


#### Retrieving all initiator groups for Linux

The `os_type` query parameter is used to perform the query.

```
# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups?os_type=linux' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup1",
    "os_type": "linux",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups?os_type=linux"
  }
}
}
```

### Retrieving a specific initiator group

```
# The API:
GET /api/protocols/san/igroups/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
    }
  }
},
"uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
"name": "igroup1",
"protocol": "iscsi",
"os_type": "linux",
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
  }
}
}
```

#### Retrieving LUNs mapped to a specific initiator group

The `fields` parameter is used to specify the desired properties.

```
# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072?fields=lun_maps' -H 'Accept: application/hal+json'

# The response:
{
```

```
"svm": {
  "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
    }
  }
},
"uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
"name": "igroup1",
"lun_maps": [
  {
    "logical_unit_number": 0,
    "lun": {
      "name": "/vol/vol1/lun1",
      "uuid": "4b33ba57-c4e0-4dbb-bc47-214800d18a71",
      "node": {
        "name": "node1",
        "uuid": "f17182af-223f-4d51-8197-2cb2146d5c4c",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/f17182af-223f-4d51-8197-
2cb2146d5c4c"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/storage/luns/4b33ba57-c4e0-4dbb-bc47-214800d18a71"
        }
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
  }
}
}
```

#### Renaming an initiator group

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

```
# The API:
PATCH /api/protocols/san/igroups/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072' -H 'Accept: application/hal+json' -d '{ "name":
"igroup1_newName" }'
```

---

#### Changing the operating system type of an initiator group

```
# The API:
PATCH /api/protocols/san/igroups/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072' -H 'Accept: application/hal+json' -d '{ "os_type":
"aix" }'
```

---

#### Adding an initiator to an initiator group

```
# The API:
POST /api/protocols/san/igroups/{igroup.uuid}/initiators

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/initiators' -H 'Accept: application/hal+json' -d '{
"name": "iqn.1991-05.com.ms:host2" }'
```

---

#### Adding multiple initiators to an initiator group

Note the use of the `records` property to add multiple initiators to the initiator group in a single API call.

```
# The API:
POST /api/protocols/san/igroups/{igroup.uuid}/initiators

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/initiators' -H 'Accept: application/hal+json' -d '{
"records": [ { "name": "iqn.1991-05.com.ms:host3" }, { "name": "iqn.1991-
05.com.ms:host4" } ] }'
```

#### Removing an initiator from an initiator group

```
# The API:
DELETE /api/protocols/san/igroups/{igroup.uuid}/initiators/iqn.1991-
05.com.ms:host3

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/initiators/iqn.1991-05.com.ms:host3' -H 'Accept:
application/hal+json'
```

#### Removing an initiator from a mapped initiator group

Normally, removing an initiator from an initiator group that is mapped to a LUN is not allowed. The removal can be forced using the `allow_delete_while_mapped` query parameter.

```
# The API:
DELETE /api/protocols/san/igroups/{igroup.uuid}/initiators/iqn.1991-
05.com.ms:host4

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/initiators/iqn.1991-
05.com.ms:host4?allow_delete_while_mapped=true' -H 'Accept:
application/hal+json'
```

#### Adding a nested initiator group to an initiator group

```
# The API:
POST /api/protocols/san/igroups/{igroup.uuid}/igroups

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/igroups' -H 'Accept: application/hal+json' -d '{
"name": "host2_igroup" }'
```

---

#### Adding multiple nested initiator groups to an initiator group

Note the use of the `records` property to add multiple nested initiator groups to the initiator group in a single API call.

```
# The API:
POST /api/protocols/san/igroups/{igroup.uuid}/igroups

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/igroups' -H 'Accept: application/hal+json' -d '{
"records": [ { "name": "host3_igroup" }, { "uuid": "c439efc8-0a70-11eb-
adc1-0242ac120002" } ] }'
```

---

#### Removing a nested initiator group from an initiator group

```
# The API:
DELETE /api/protocols/san/igroups/{igroup.uuid}/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/igroups/c439efc8-0a70-11eb-adc1-0242ac120002' -H
'Accept: application/hal+json'
```

---

#### Removing a nested initiator group from a mapped initiator group

Normally, removing a nested initiator group from an initiator group that is mapped to a LUN is not allowed. The removal can be forced using the `allow_delete_while_mapped` query parameter.

```
# The API:
DELETE /api/protocols/san/igroups/{igroup.uuid}/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/igroups/c439efc8-0a70-11eb-adc1-
0242ac120002?allow_delete_while_mapped=true' -H 'Accept:
application/hal+json'
```

---

#### Deleting an initiator group

```
# The API:
DELETE /api/protocols/san/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/abf9c39d-ab9f-
11e8-b8a3-005056bb7072' -H 'Accept: application/hal+json'
```

---

#### Deleting a mapped initiator group

Normally, deleting an initiator group that is mapped to a LUN is not allowed. The deletion can be forced using the `allow_delete_while_mapped` query parameter.

```
# The API:
DELETE /api/protocols/san/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/abf9c39d-ab9f-
11e8-b8a3-005056bb7072?allow_delete_while_mapped=true' -H 'Accept:
application/hal+json'
```

#### Retrieve initiator groups

GET `/protocols/san/igroups`

**Introduced In:** 9.6

Retrieves initiator groups.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connectivity_tracking.*`
- `igroups.*`
- `lun_maps.*`
- `parent_igroups.*`
- `target.*`

#### Related ONTAP commands

- `lun igroup show`
- `lun mapping show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| target.vendor_id | string | query | False | Filter by target.vendor_id<br>• Introduced in: 9.11 |
| target.product_id | string | query | False | Filter by target.product_id<br>• Introduced in: 9.11 |
| target.firmware_revision | string | query | False | Filter by target.firmware_revision<br>• Introduced in: 9.11 |
| supports_igroups | boolean | query | False | Filter by supports_igroups<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| name | string | query | False | Filter by name<br>• maxLength: 96<br>• minLength: 1 |
| replication.state | string | query | False | Filter by replication.state<br>• Introduced in: 9.15 |
| replication.error.summary.arguments.message | string | query | False | Filter by replication.error.summary.arguments.message<br>• Introduced in: 9.15 |
| replication.error.summary.arguments.code | string | query | False | Filter by replication.error.summary.arguments.code<br>• Introduced in: 9.15 |
| replication.error.summary.code | string | query | False | Filter by replication.error.summary.code<br>• Introduced in: 9.15 |
| replication.error.summary.message | string | query | False | Filter by replication.error.summary.message<br>• Introduced in: 9.15 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| replication.error.igroup.name | string | query | False | Filter by replication.error.igroup.name<br>• Introduced in: 9.15<br>• maxLength: 96<br>• minLength: 1 |
| replication.error.igroup.local_svm | boolean | query | False | Filter by replication.error.igroup.local_svm<br>• Introduced in: 9.15 |
| replication.error.igroup.uuid | string | query | False | Filter by replication.error.igroup.uuid<br>• Introduced in: 9.15 |
| replication.peer_svm.name | string | query | False | Filter by replication.peer_svm.name<br>• Introduced in: 9.15 |
| replication.peer_svm.uuid | string | query | False | Filter by replication.peer_svm.uuid<br>• Introduced in: 9.15 |
| igroups.uuid | string | query | False | Filter by igroups.uuid<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| igroups.name | string | query | False | Filter by igroups.name<br>• Introduced in: 9.9<br>• maxLength: 96<br>• minLength: 1 |
| igroups.comment | string | query | False | Filter by igroups.comment<br>• Introduced in: 9.9<br>• maxLength: 254<br>• minLength: 0 |
| protocol | string | query | False | Filter by protocol |
| uuid | string | query | False | Filter by uuid |
| comment | string | query | False | Filter by comment<br>• Introduced in: 9.9<br>• maxLength: 254<br>• minLength: 0 |
| connectivity_tracking.connection_state | string | query | False | Filter by connectivity_tracking.connection_state<br>• Introduced in: 9.11 |
| connectivity_tracking.required_nodes.name | string | query | False | Filter by connectivity_tracking.required_nodes.name<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| connectivity_tracking.required_nodes.uuid | string | query | False | Filter by connectivity_tracking.required_nodes.uuid<br><br>• Introduced in: 9.11 |
| connectivity_tracking.alerts.summary.arguments.message | string | query | False | Filter by connectivity_tracking.alerts.summary.arguments.message<br><br>• Introduced in: 9.11 |
| connectivity_tracking.alerts.summary.arguments.code | string | query | False | Filter by connectivity_tracking.alerts.summary.arguments.code<br><br>• Introduced in: 9.11 |
| connectivity_tracking.alerts.summary.code | string | query | False | Filter by connectivity_tracking.alerts.summary.code<br><br>• Introduced in: 9.11 |
| connectivity_tracking.alerts.summary.message | string | query | False | Filter by connectivity_tracking.alerts.summary.message<br><br>• Introduced in: 9.11 |
| initiators.proximity.peer_svms.name | string | query | False | Filter by initiators.proximity.peer_svms.name<br><br>• Introduced in: 9.15 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| initiators.proximity.peer_svms.uuid | string | query | False | Filter by initiators.proximity.peer_svms.uuid<br>• Introduced in: 9.15 |
| initiators.proximity.local_svm | boolean | query | False | Filter by initiators.proximity.local_svm<br>• Introduced in: 9.15 |
| initiators.name | string | query | False | Filter by initiators.name<br>• maxLength: 96<br>• minLength: 1 |
| initiators.comment | string | query | False | Filter by initiators.comment<br>• Introduced in: 9.9<br>• maxLength: 254<br>• minLength: 0 |
| initiators.connectivity_tracking.connection_state | string | query | False | Filter by initiators.connectivity_tracking.connection_state<br>• Introduced in: 9.11 |
| initiators.igroup.name | string | query | False | Filter by initiators.igroup.name<br>• Introduced in: 9.9<br>• maxLength: 96<br>• minLength: 1 |
| initiators.igroup.uuid | string | query | False | Filter by initiators.igroup.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| delete_on_unmap | boolean | query | False | Filter by delete_on_unmap |
| portset.uuid | string | query | False | Filter by portset.uuid<br>• Introduced in: 9.9 |
| portset.name | string | query | False | Filter by portset.name<br>• Introduced in: 9.9<br>• maxLength: 96<br>• minLength: 1 |
| lun_maps.logical_unit_number | integer | query | False | Filter by lun_maps.logical_unit_number |
| lun_maps.lun.name | string | query | False | Filter by lun_maps.lun.name |
| lun_maps.lun.node.name | string | query | False | Filter by lun_maps.lun.node.name |
| lun_maps.lun.node.uuid | string | query | False | Filter by lun_maps.lun.node.uuid |
| lun_maps.lun.uuid | string | query | False | Filter by lun_maps.lun.uuid |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| os_type | string | query | False | Filter by os_type |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| parent_igroups.comment | string | query | False | Filter by parent_igroups.comment<br>• Introduced in: 9.9<br>• maxLength: 254<br>• minLength: 0 |
| parent_igroups.name | string | query | False | Filter by parent_igroups.name<br>• Introduced in: 9.9<br>• maxLength: 96<br>• minLength: 1 |
| parent_igroups.uuid | string | query | False | Filter by parent_igroups.uuid<br>• Introduced in: 9.9 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[igroup] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "connectivity_tracking": {
        "alerts": [
          {
            "summary": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
              ],
              "code": "4",
              "message": "entry doesn't exist"
            }
          }
        ],
        "connection_state": "string",
        "required_nodes": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
```

```
        ]
      },
      "igroups": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "comment": "string",
          "igroups": [
            null
          ],
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "initiators": [
        {
          "_links": {
            "connectivity_tracking": {
              "href": "/api/resourcelink"
            },
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "comment": "string",
          "connectivity_tracking": {
            "connection_state": "string"
          },
          "igroup": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "name": "iqn.1998-01.com.corp.iscsi:name1",
          "proximity": {
            "peer_svms": [
              {
                "_links": {
                  "self": {
```

```
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
  ],
  "lun_maps": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "logical_unit_number": 0,
      "lun": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "lun1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "parent_igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
          },
          "comment": "string",
          "name": "igroup1",
          "parent_igroups": [
            null
          ],
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "portset": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "portset1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "protocol": "string",
      "replication": {
        "error": {
          "igroup": {
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "summary": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          }
        },
        "peer_svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "state": "string"
```

```
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "firmware_revision": "9111",
        "product_id": "LUN C-Mode",
        "vendor_id": "NETAPP"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state of the initiator group.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

| Name | Type | Description |
|---|---|---|
| summary | summary | A user friendly message describing the connection state of the initiator group. |

required_nodes

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

connectivity_tracking

An overview of the connections to ONTAP by the initiators in this group.

| Name | Type | Description |
| --- | --- | --- |
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| required_nodes | array[required_nodes] | Nodes to which the initiators in this group should be connected to ensure reliable service. This is the collection of any node hosting a LUN mapped to this igroup as well as the HA partners of those nodes. |

igroup_child

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroup_child] | Further nested initiator groups. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

connectivity_tracking

A link to the initiator with connectivity information relevant to its membership of this initiator group.

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

self

A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group,

the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group.

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| connectivity_tracking | connectivity_tracking | A link to the initiator with connectivity information relevant to its membership of this initiator group. |
| self | self | A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group, the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group. |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
|---|---|---|
| connection_state | string | Connection state. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

initiators

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

lun

The LUN to which the initiator group is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the LUN. |
| node | node | |
| uuid | string | The unique identifier of the LUN. |

lun_maps

A LUN map with which the initiator group is associated.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| logical_unit_number | integer | The logical unit number assigned to the LUN for initiators in the initiator group. |
| lun | lun | The LUN to which the initiator group is mapped. |

igroup_parent

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment available for use by the administrator. |
| name | string | The name of the initiator group. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| uuid | string | The unique identifier of the initiator group. |

portset

The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.

In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.

Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string ("").

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the portset. |
| uuid | string | The unique identifier of the portset. |

igroup

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | Indicates whether the reported igroup is on the local SVM or the peer SVM. When deleting a replicated igroup, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer igroup remains, the peer igroup is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

summary

A user friendly message describing the error.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
|---|---|---|
| igroup | igroup | |
| summary | summary | A user friendly message describing the error. |

peer_svm

The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to initiator group replication.

| Name | Type | Description |
|---|---|---|
| error | error | Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |

| Name | Type | Description |
| --- | --- | --- |
| peer_svm | peer_svm | The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`. |
| state | string | The state of the replication queue associated with this igroup. If this igroup is not in the replication queue, the state is reported as *ok*. If this igroup is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this igroup is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

Properties of the SCSI target to which the initiator group provides access.

| Name | Type | Description |
| --- | --- | --- |
| firmware_revision | string | The firmware revision of the SCSI target specific to the OS type of the initiator group. |

| Name | Type | Description |
|---|---|---|
| product_id | string | The product ID of the SCSI target. |
| vendor_id | string | The vendor ID of the SCSI target. |

igroup

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPNs), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host endpoints.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the host initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator group may contain either initiators or other initiator groups, but not both simultaneously. When a parent initiator group is mapped, it inherits all of the initiators of any initiator groups nested below it. If any nested initiator group is modified to contain different initiators, the parent initiator groups inherit the change. A parent can have many nested initiator groups and an initiator group can be nested under multiple parents. Initiators can only be added or removed from the initiator group that directly contains them. The maximum supported depth of nesting is three layers.

Best practice when using nested initiator groups is to match host hierarchies. A single initiator group should correspond to a single host. If a LUN needs to be mapped to multiple hosts, the initiator groups representing those hosts should be aggregated into a parent initiator group and the LUN should be mapped to that initiator group. For multi-ported hosts, initiators have a comment property where the port corresponding to the initiator can be documented.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once. With the introduction of nestable initiator groups, best practice is to use the hierarchy such that an initiator is only a direct member of a single initiator group, and that initiator group can then be referenced by other initiator groups.

All initiators or nested initiator groups in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (*fcp*), iSCSI (*iscsi*), or both (*mixed*). Initiator groups within a nested hierarchy may not have conflicting protocols.

Zero or more initiators or nested initiator groups can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. Initiator groups containing other initiator groups report the aggregated list of initiators from all nested initiator groups, but modifications of the initiator list must be performed on the initiator group that directly contains the initiators. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroup_child] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[initiators] | The initiators that are members of the group or any group nested below this group. Optional in POST.<br><br>This property is mutually exclusive with the *igroups* property during POST.<br><br>During GET, this array contains initiators that are members of this group or any nested initiator groups below this group. When initiators of nested groups are returned, they include links to the initiator group that directly contains the initiator.<br><br>Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details. |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| replication | replication | Properties related to initiator group replication. |

| Name | Type | Description |
|---|---|---|
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an initiator group**

POST `/protocols/san/igroups`

**Introduced In:** 9.6

Creates an initiator group.

**Required properties**

- `svm.uuid` or `svm.name` - Existing SVM in which to create the initiator group.

- `name` - Name of the initiator group.
- `os_type` - Operating system of the initiator group's initiators.

#### Recommended optional properties

- `initiators.name` - Name(s) of initiator group's initiators. This property can be used to create the initiator group and populate it with initiators in a single request.

#### Default property values

If not specified in POST, the following default property values are assigned.

- `protocol` - *mixed* - Data protocol of the initiator group's initiators.

#### Related ONTAP commands

- `lun igroup create`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |

| Name | Type | Description |
| --- | --- | --- |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |
| igroups | array[igroup_child] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[initiators] | The initiators that are members of the group or any group nested below this group. Optional in POST.<br><br>This property is mutually exclusive with the *igroups* property during POST.<br><br>During GET, this array contains initiators that are members of this group or any nested initiator groups below this group. When initiators of nested groups are returned, they include links to the initiator group that directly contains the initiator.<br><br>Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details. |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| replication | replication | Properties related to initiator group replication. |

| Name | Type | Description |
| --- | --- | --- |
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

**Example request**


```
{
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {}
    ],
    "connection_state": "string",
    "required_nodes": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  },
  "igroups": [
    {
      "comment": "string",
      "igroups": [
        null
      ],
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiators": [
    {
      "comment": "string",
      "connectivity_tracking": {
        "connection_state": "string"
      },
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
```

```
  ],
  "lun_maps": [
    {
      "logical_unit_number": 0,
      "lun": {
        "name": "lun1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "parent_igroups": [
    {
      "comment": "string",
      "name": "igroup1",
      "parent_igroups": [
        null
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "portset": {
    "name": "portset1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "protocol": "string",
  "replication": {
    "error": {
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "peer_svm": {
      "name": "peer1",
      "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
    },
    "state": "string"
  },
  "svm": {
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "firmware_revision": "9111",
    "product_id": "LUN C-Mode",
    "vendor_id": "NETAPP"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[igroup] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "comment": "string",
      "connectivity_tracking": {
        "alerts": [
          {}
        ],
        "connection_state": "string",
        "required_nodes": [
          {
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ]
      },
      "igroups": [
        {
          "comment": "string",
          "igroups": [
            null
          ],
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "initiators": [
        {
          "comment": "string",
          "connectivity_tracking": {
            "connection_state": "string"
          },
          "igroup": {
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "name": "iqn.1998-01.com.corp.iscsi:name1",
          "proximity": {
            "peer_svms": [
              {
                "name": "peer1",
                "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
              }
```

```
            ]
          }
        }
      ],
      "lun_maps": [
        {
          "logical_unit_number": 0,
          "lun": {
            "name": "lun1",
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "name": "igroup1",
      "os_type": "string",
      "parent_igroups": [
        {
          "comment": "string",
          "name": "igroup1",
          "parent_igroups": [
            null
          ],
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "portset": {
        "name": "portset1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "protocol": "string",
      "replication": {
        "error": {
          "igroup": {
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        },
        "peer_svm": {
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "state": "string"
```

```
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "firmware_revision": "9111",
        "product_id": "LUN C-Mode",
        "vendor_id": "NETAPP"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The supplied SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373958 | An invalid initiator group name was supplied. |
| 5373966 | An initiator group cannot be created in an SVM that is configured for NVMe. |
| 5373969 | A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing. |
| 5373971 | A supplied initiator name looks like an iSCSI IQN initiator, but the date portion is invalid. |
| 5373972 | A supplied initiator name looks like an iSCSI IQN initiator, but the naming authority portion is invalid. |

| Error Code | Description |
| --- | --- |
| 5373977 | A supplied initiator name looks like an iSCSI EUI initiator, but the length is invalid. |
| 5373978 | A supplied initiator name looks like an iSCSI EUI initiator, but the format is invalid. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5373992 | A supplied initiator name was too long to be valid. |
| 5373993 | A supplied initiator name did not match any valid format. |
| 5374023 | An initiator group with the same name already exists. |
| 5374027 | An attempt was made to bind a portset with no member network interfaces to the initiator group. |
| 5374028 | An attempt was made to bind a portset with an incompatible protocol to the initiator group. |
| 5374038 | An invalid Fibre Channel WWPN was supplied. |
| 5374039 | An invalid iSCSI initiator name was supplied. |
| 5374040 | Initiators and child initiator groups were both supplied, but only one option is allowed. |
| 5374732 | An initiator is already in another initiator group with a conflicting operating system type. |
| 5374735 | An attempt was made to add a child igroup that would exceed the maximum allowable depth. |
| 5374737 | A supplied child initiator group already exists in another initiator group's hierarchy. |
| 5374739 | A supplied child initiator group has an operating system type that differs from the parent initiator group. |
| 5374740 | A supplied child initiator group has an protocol that differs from the parent initiator group. |
| 5374741 | A supplied child initiator group is already owned by a different child in the initiator group's hierarchy. |
| 5374742 | A supplied child initiator group contains an initiator that is already owned by another initiator group in the hierarchy. |
| 5374745 | Initiator group cannot be added as a child to itself. |
| 5374746 | The cluster is currently running in a mixed version and nested initiator groups cannot be created until the effective cluster version reaches 9.9.1. |

| Error Code | Description |
| --- | --- |
| 5374747 | The cluster is currently running in a mixed version and initiator group comments cannot be created until the effective cluster version reaches 9.9.1. |
| 5374758 | An error was reported by the peer cluster while creating a replicated initiator group. The specific error will be included as a nested error. |
| 5374878 | The supplied child initiator group does not exist. |
| 5374911 | The supplied portset does not exist. |
| 5374917 | Duplicated initiators have conflicting property values. |
| 5375055 | The `local_svm` property of an initiator proximity was not specified. |
| 5375056 | An SVM peering relationship that does not have the initiator group's SVM as the local SVM was specified. |
| 5375258 | The igroup is already replicated to a different peer SVM. |
| 5375261 | Setting initiator proximity is not supported for the SVM type. |
| 5376057 | Setting initiator proximity is not supported for the ONTAP version. |
| 5376059 | Setting initiator proximity to a peer that is either the destination of an SVM DR relationship or in a Metrocluster configuration is not supported. |
| 5376253 | Initiator group replication requires an effective cluster version of 9.15.1. |
| 5376255 | Initiator group replication requires the peer cluster to have an effective cluster version of 9.15.1. |
| 5376488 | An NVMe over Fabrics subsystem already exists with the requested name. |
| 6620376 | SVM peering information is unavailable. |
| 6620384 | The supplied SVMs are not peered. |
| 26345672 | The specified SVM peering relationship was not found. |
| 26345673 | An SVM peering relationship between the initiator group's SVM and specified peer SVM was not found. |
| 26345675 | An SVM peering relationship UUID and name were specified and they do not refer to the same SVM peering relationship. |
| 26345680 | Supplied SVM peer is on the local cluster. The operation requires a peer on a remote cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state of the initiator group.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

required_nodes

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

connectivity_tracking

An overview of the connections to ONTAP by the initiators in this group.

| Name | Type | Description |
|---|---|---|
| alerts | array[alerts] | |
| connection_state | string | Connection state. |

| Name | Type | Description |
| --- | --- | --- |
| required_nodes | array[required_nodes] | Nodes to which the initiators in this group should be connected to ensure reliable service. This is the collection of any node hosting a LUN mapped to this igroup as well as the HA partners of those nodes. |

igroup_child

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroup_child] | Further nested initiator groups. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

connectivity_tracking

A link to the initiator with connectivity information relevant to its membership of this initiator group.

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

self

A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group, the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| connection_state | string | Connection state. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

initiators

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

lun

The LUN to which the initiator group is mapped.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. |
| node | node | |
| uuid | string | The unique identifier of the LUN. |

lun_maps

A LUN map with which the initiator group is associated.

| Name | Type | Description |
|---|---|---|
| logical_unit_number | integer | The logical unit number assigned to the LUN for initiators in the initiator group. |
| lun | lun | The LUN to which the initiator group is mapped. |

igroup_parent

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| name | string | The name of the initiator group. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| uuid | string | The unique identifier of the initiator group. |

portset

The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.

In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.

Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string ("").

| Name | Type | Description |
|---|---|---|
| name | string | The name of the portset. |
| uuid | string | The unique identifier of the portset. |

igroup

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | Indicates whether the reported igroup is on the local SVM or the peer SVM. When deleting a replicated igroup, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer igroup remains, the peer igroup is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

summary

A user friendly message describing the error.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | |

peer_svm

The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to initiator group replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_svm | peer_svm | The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | The state of the replication queue associated with this igroup. If this igroup is not in the replication queue, the state is reported as *ok*. If this igroup is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this igroup is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

Properties of the SCSI target to which the initiator group provides access.

| Name | Type | Description |
| --- | --- | --- |
| firmware_revision | string | The firmware revision of the SCSI target specific to the OS type of the initiator group. |
| product_id | string | The product ID of the SCSI target. |
| vendor_id | string | The vendor ID of the SCSI target. |

igroup

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPNs), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host

endpoints.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the host initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator group may contain either initiators or other initiator groups, but not both simultaneously. When a parent initiator group is mapped, it inherits all of the initiators of any initiator groups nested below it. If any nested initiator group is modified to contain different initiators, the parent initiator groups inherit the change. A parent can have many nested initiator groups and an initiator group can be nested under multiple parents. Initiators can only be added or removed from the initiator group that directly contains them. The maximum supported depth of nesting is three layers.

Best practice when using nested initiator groups is to match host hierarchies. A single initiator group should correspond to a single host. If a LUN needs to be mapped to multiple hosts, the initiator groups representing those hosts should be aggregated into a parent initiator group and the LUN should be mapped to that initiator group. For multi-ported hosts, initiators have a comment property where the port corresponding to the initiator can be documented.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once. With the introduction of nestable initiator groups, best practice is to use the hierarchy such that an initiator is only a direct member of a single initiator group, and that initiator group can then be referenced by other initiator groups.

All initiators or nested initiator groups in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (*fcp*), iSCSI (*iscsi*), or both (*mixed*). Initiator groups within a nested hierarchy may not have conflicting protocols.

Zero or more initiators or nested initiator groups can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. Initiator groups containing other initiator groups report the aggregated list of initiators from all nested initiator groups, but modifications of the initiator list must be performed on the initiator group that directly contains the initiators. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |

| Name | Type | Description |
| --- | --- | --- |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |
| igroups | array[igroup_child] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[initiators] | The initiators that are members of the group or any group nested below this group. Optional in POST.<br><br>This property is mutually exclusive with the *igroups* property during POST.<br><br>During GET, this array contains initiators that are members of this group or any nested initiator groups below this group. When initiators of nested groups are returned, they include links to the initiator group that directly contains the initiator.<br><br>Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details. |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| replication | replication | Properties related to initiator group replication. |

| Name | Type | Description |
| --- | --- | --- |
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve nested initiator groups of an initiator group**

GET `/protocols/san/igroups/{igroup.uuid}/igroups`

**Introduced In:** 9.9

Retrieves nested initiator groups of an initiator group.

This API only reports the nested initiator groups that are direct children of the initiator group. Further nested initiator groups are reported by their direct parent initiator group.

#### Related ONTAP commands

- `lun igroup show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the parent initiator group. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[igroup_nested] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374852 | The initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the parent initiator group. |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

igroup_nested

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | |
| name | string | The name of the initiator group. |

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the initiator group. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add nested initiator groups to an initiator group

POST `/protocols/san/igroups/{igroup.uuid}/igroups`

**Introduced In:** 9.9

Adds one or more nested initiator groups to an initiator group. A single nested initiator group can be added by directly specifying the name or UUID. Multiple nested initiator groups can be added by specifying the names or UUIDs in the records array. Nested initiator groups cannot be added to an initiator group that already directly contains initiators.

#### Required properties

- `name` and/or `uuid` or `records` - Nested initiator groups to add to the initiator group.

#### Related ONTAP commands

- `lun igroup add`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the parent initiator group. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| igroup | igroup | |
| name | string | The name of the initiator group. |
| records | array[records] | An array of initiator groups specified to add multiple nested initiator groups to an initiator group in a single API call. Not allowed when the `name` property is used. |
| uuid | string | The unique identifier of the initiator group. |

#### Example request

```
{
  "igroup": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "igroup1",
  "records": [
    {
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[igroup_nested] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "igroup": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "igroup1",
      "records": [
        {
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5374735 | An attempt was made to add a child igroup that would exceed the maximum allowable depth. |
| 5374736 | A supplied child initiator group already exists in the parent initiator group's hierarchy. |
| 5374737 | A supplied child initiator group already exists in another initiator group's hierarchy. |
| 5374739 | A supplied child initiator group has an operating system type that differs from the parent initiator group. |
| 5374740 | A supplied child initiator group has an protocol that differs from the parent initiator group. |
| 5374741 | A supplied child initiator group is already owned by a different child in the initiator group's hierarchy. |
| 5374742 | A supplied child initiator group contains an initiator that is already owned by another initiator group in the hierarchy. |
| 5374764 | An unreplicated initiator group cannot be added as a child of a replicated initiator group. |
| 5374852 | The initiator group specified in the URI does not exist. |
| 5374878 | The supplied child initiator group does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

igroup

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the parent initiator group. |

records

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

igroup_nested

| Name | Type | Description |
|---|---|---|
| igroup | igroup | |
| name | string | The name of the initiator group. |
| records | array[records] | An array of initiator groups specified to add multiple nested initiator groups to an initiator group in a single API call. Not allowed when the `name` property is used. |
| uuid | string | The unique identifier of the initiator group. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
|---|---|---|
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Remove a nested initiator group from an initiator group**

DELETE `/protocols/san/igroups/{igroup.uuid}/igroups/{uuid}`

**Introduced In:** 9.9

Removes a nested initiator group from an initiator group. This API does not delete the nested initiator group itself. It removes the relationship between a parent and child initiator group.

This API only supports removal of initiator groups owned directly by the initiator group. Further nested initiator groups must be removed from the direct parent initiator group.

**Related ONTAP commands**

- `lun igroup remove`

**Learn more**

- DOC /protocols/san/igroups

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the parent initiator group. |
| uuid | string | path | True | The unique identifier of the nested initiator group. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| allow_delete_while_mapped | boolean | query | False | Allows the deletion of a nested initiator group from of a mapped initiator group.<br><br>Deleting a nested initiator group from a mapped initiator group means that the LUNs, to which the initiator group is mapped, are no longer available to the initiators nested below the initiator group being removed. This might cause a disruption in the availability of data.<br><br>**This parameter should be used with caution.**<br><br>• Default value: |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254213 | The initiator group is mapped to one or more LUNs and `allow_delete_while_mapped` has not been specified. |
| 5374738 | The child initiator group is not owned by the parent initiator group. |
| 5374743 | LUN maps exist for a parent initiator group. |

| Error Code | Description |
| --- | --- |
| 5374852 | The initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a nested initiator group of an initiator group

GET `/protocols/san/igroups/{igroup.uuid}/igroups/{uuid}`

**Introduced In:** 9.9

Retrieves a nested initiator group of an initiator group.

This API only reports the nested initiator groups that are direct children of the initiator group. Further nested initiator groups are reported by their direct parent initiator group.

#### Related ONTAP commands

- `lun igroup show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the parent initiator group. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the nested initiator group. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "igroup1",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The nested initiator group is not a member of the initiator group. |
| 5374852 | The parent initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the parent initiator group. |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve initiators of an initiator group

GET `/protocols/san/igroups/{igroup.uuid}/initiators`

**Introduced In:** 9.6

Retrieves initiators of an initiator group.

This API only reports initiators owned directly by the initiator group. Initiators of nested initiator groups are not included in this collection.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connectivity_tracking.*`

#### Related ONTAP commands

- `lun igroup show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the initiator group. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[igroup_initiator] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "connectivity_tracking": {
        "alerts": [
          {
            "summary": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
              ],
              "code": "4",
              "message": "entry doesn't exist"
            }
          }
        ],
        "connection_state": "string",
        "connections": [
          {
            "logins": [
              {
                "interface": {
                  "fc": {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
```

```
                    },
                    "name": "fc_lif1",
                    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                    "wwpn": "20:00:00:50:56:b4:13:a8"
                  },
                  "ip": {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "ip": {
                      "address": "10.10.10.7"
                    },
                    "name": "lif1",
                    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                  }
                },
                "last_seen_time": "2021-03-14 05:19:00 +0000"
              }
            ],
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "proximity": {
        "peer_svms": [
          {
```

```
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374852 | The initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

| Name | Type | Description |
|---|---|---|
| summary | summary | A user friendly message describing the connection state. |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interface

| Name | Type | Description |
|---|---|---|
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |

logins

| Name | Type | Description |
|---|---|---|
| connected | boolean | True if the initiator is currently logged in to this connection's interface. |
| interface | interface | |
| last_seen_time | string | The last time this initiator logged in. Logins not seen for 48 hours are cleared and not reported. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

connections

| Name | Type | Description |
|---|---|---|
| logins | array[logins] | |
| node | node | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
|---|---|---|
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| connections | array[connections] | |

igroup

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

records

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

igroup_initiator

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br><br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Add initiators to an initiator group**

POST `/protocols/san/igroups/{igroup.uuid}/initiators`

**Introduced In:** 9.6

Adds one or more initiators to an initiator group.

This API does not support adding initiators to an initiator group that already contains nested initiator groups.

#### Required properties

- `name` or `records.name` - Initiator name(s) to add to the initiator group.

#### Related ONTAP commands

- `lun igroup add`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the initiator group. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |
| records | array[records] | An array of initiators specified to add multiple initiators to an initiator group in a single API call. Not allowed when the `name` property is used. |

**Example request**


```
{
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {}
    ],
    "connection_state": "string",
    "connections": [
      {
        "logins": [
          {
            "interface": {
              "fc": {
                "name": "fc_lif1",
                "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                "wwpn": "20:00:00:50:56:b4:13:a8"
              },
              "ip": {
                "name": "lif1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            },
            "last_seen_time": "2021-03-14 05:19:00 +0000"
          }
        ],
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "igroup": {
    "name": "igroup1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "iqn.1998-01.com.corp.iscsi:name1",
  "proximity": {
    "peer_svms": [
      {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
```

```
  },
  "records": [
    {
      "comment": "string",
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
  ]
}
```


**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[igroup_initiator] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "comment": "string",
      "connectivity_tracking": {
        "alerts": [
          {}
        ],
        "connection_state": "string",
        "connections": [
          {
            "logins": [
              {
                "interface": {
                  "fc": {
                    "name": "fc_lif1",
                    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                    "wwpn": "20:00:00:50:56:b4:13:a8"
                  },
                  "ip": {
                    "name": "lif1",
                    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                  }
                },
                "last_seen_time": "2021-03-14 05:19:00 +0000"
              }
            ],
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      },
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
```

```
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      },
      "records": [
        {
          "comment": "string",
          "name": "iqn.1998-01.com.corp.iscsi:name1",
          "proximity": {
            "peer_svms": [
              {
                "name": "peer1",
                "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
              }
            ]
          }
        }
      ]
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254193 | Adding an initiator would cause the initiator to be mapped to the same LUN more than once. |
| 1254324 | Adding an initiator would cause the initiator to have the same logical unit identifier for multiple LUN maps. |
| 5373969 | A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing. |
| 5373971 | A supplied initiator name looks like an iSCSI IQN initiator, but the date portion is invalid. |

| Error Code | Description |
| --- | --- |
| 5373972 | A supplied initiator name looks like an iSCSI IQN initiator, but the naming authority portion is invalid. |
| 5373977 | A supplied initiator name looks like an iSCSI EUI initiator, but the length is invalid. |
| 5373978 | A supplied initiator name looks like an iSCSI EUI initiator, but the format is invalid. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5373992 | A supplied initiator name was too long to be valid. |
| 5373993 | A supplied initiator name did not match any valid format. |
| 5374033 | Initiators must be supplied. |
| 5374035 | A supplied initiator is already in the initiator group. |
| 5374038 | An invalid Fibre Channel WWPN was supplied. |
| 5374039 | An invalid iSCSI initiator name was supplied. |
| 5374734 | An initiator is already in another initiator group with a conflicting operating system type. |
| 5374761 | An error was reported by the peer cluster while adding an initiator to a replicated initiator group. The specific error will be included as a nested error. |
| 5374852 | The initiator group specified in the URI does not exist. |
| 5374917 | Duplicated initiators have conflicting property values. |
| 5375055 | The `local_svm` property of an initiator proximity was not specified. |
| 5375056 | An SVM peering relationship that does not have the initiator group's SVM as the local SVM was specified. |
| 5375261 | Setting initiator proximity is not supported for the SVM type. |
| 5376057 | Setting initiator proximity is not supported for the ONTAP version. |
| 5376059 | Setting initiator proximity to a peer that is either the destination of an SVM DR relationship or in a Metrocluster configuration is not supported. |
| 26345672 | The specified SVM peering relationship was not found. |
| 26345673 | An SVM peering relationship between the initiator group's SVM and specified peer SVM was not found. |

| Error Code | Description |
| --- | --- |
| 26345675 | An SVM peering relationship UUID and name were specified and they do not refer to the same SVM peering relationship. |
| 26345680 | Supplied SVM peer is on the local cluster. The operation requires a peer on a remote cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interface

| Name | Type | Description |
|---|---|---|
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |

logins

| Name | Type | Description |
|---|---|---|
| connected | boolean | True if the initiator is currently logged in to this connection's interface. |
| interface | interface | |
| last_seen_time | string | The last time this initiator logged in. Logins not seen for 48 hours are cleared and not reported. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

connections

| Name | Type | Description |
|---|---|---|
| logins | array[logins] | |
| node | node | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
|---|---|---|
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| connections | array[connections] | |

igroup

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator

in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

records

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

igroup_initiator

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br><br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |
| records | array[records] | An array of initiators specified to add multiple initiators to an initiator group in a single API call. Not allowed when the `name` property is used. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an initiator from an initiator group

DELETE `/protocols/san/igroups/{igroup.uuid}/initiators/{name}`

**Introduced In:** 9.6

Deletes an initiator from an initiator group.

This API only supports removal of initiators owned directly by the initiator group. Initiators of nested initiator groups must be removed on the initiator group that directly owns the initiator.

#### Related ONTAP commands

- `lun igroup remove`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the initiator group. |
| name | string | path | True | The initiator name. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| allow_delete_while_mapped | boolean | query | False | Allows the deletion of an initiator from of a mapped initiator group.<br><br>Deleting an initiator from a mapped initiator group makes the LUNs to which the initiator group is mapped no longer available to the initiator. This might cause a disruption in the availability of data.<br><br>**This parameter should be used with caution.**<br><br>• Default value: |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254213 | The initiator group is mapped to one or more LUNs and `allow_delete_while_mapped` has not been specified. |
| 5374034 | An initiator is not a member of the initiator group. |
| 5374041 | The initiator is not owned by the supplied initiator group. |
| 5374762 | An error was reported by the peer cluster while removing an initiator from a replicated initiator group. The specific error will be included as a nested error. |
| 5374852 | The initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an initiator

GET `/protocols/san/igroups/{igroup.uuid}/initiators/{name}`

**Introduced In:** 9.6

Retrieves an initiator of an initiator group.

This API only reports initiators owned directly by the initiator group. Initiators of nested initiator groups are not part of this collection.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connectivity_tracking.*`

#### Related ONTAP commands

- `lun igroup show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the initiator group. |
| name | string | path | True | The initiator name. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br><br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {
        "summary": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        }
      }
    ],
    "connection_state": "string",
    "connections": [
      {
        "logins": [
          {
            "interface": {
              "fc": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "fc_lif1",
                "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                "wwpn": "20:00:00:50:56:b4:13:a8"
              },
              "ip": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
```

```
                "ip": {
                  "address": "10.10.10.7"
                },
                "name": "lif1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            },
            "last_seen_time": "2021-03-14 05:19:00 +0000"
          }
        ],
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "igroup1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "iqn.1998-01.com.corp.iscsi:name1",
  "proximity": {
    "peer_svms": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
  }
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The initiator is not a member of the initiator group. |
| 5374852 | The initiator group specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

| Name | Type | Description |
|---|---|---|
| summary | summary | A user friendly message describing the connection state. |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interface

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |

logins

| Name | Type | Description |
| --- | --- | --- |
| connected | boolean | True if the initiator is currently logged in to this connection's interface. |
| interface | interface | |
| last_seen_time | string | The last time this initiator logged in. Logins not seen for 48 hours are cleared and not reported. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

connections

| Name | Type | Description |
| --- | --- | --- |
| logins | array[logins] | |
| node | node | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| connections | array[connections] | |

igroup

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the initiator group. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

records

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br><br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an initiator

PATCH `/protocols/san/igroups/{igroup.uuid}/initiators/{name}`

**Introduced In:** 9.9

Updates an initiator of an initiator group.

This API only supports modification of initiators owned directly by the initiator group. Initiators of nested initiator groups must be modified on the initiator group that directly owns the initiator.

#### Related ONTAP commands

- `lun igroup initiator modify`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the initiator group. |
| name | string | path | True | The initiator name. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br><br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |
| records | array[records] | An array of initiators specified to add multiple initiators to an initiator group in a single API call. Not allowed when the `name` property is used. |

**Example request**


```
{
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {}
    ],
    "connection_state": "string",
    "connections": [
      {
        "logins": [
          {
            "interface": {
              "fc": {
                "name": "fc_lif1",
                "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
                "wwpn": "20:00:00:50:56:b4:13:a8"
              },
              "ip": {
                "name": "lif1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            },
            "last_seen_time": "2021-03-14 05:19:00 +0000"
          }
        ],
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "igroup": {
    "name": "igroup1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "proximity": {
    "peer_svms": [
      {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      }
    ]
  },
```

```
  "records": [
    {
      "comment": "string",
      "proximity": {
        "peer_svms": [
          {
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
  ]
}
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374034 | An initiator is not a member of the initiator group. |
| 5374744 | The cluster is currently running in a mixed version and the initiators cannot be modified until the effective cluster version reaches 9.9.1. |
| 5374852 | The initiator group does not exist. |
| 5374918 | A subset of the provided list of initiators were modified before a failure occurred. |
| 5375055 | The `local_svm` property of an initiator proximity was not specified. |
| 5375056 | An SVM peering relationship that does not have the initiator group's SVM as the local SVM was specified. |
| 5375261 | Setting initiator proximity is not supported for the SVM type. |
| 5376057 | Setting initiator proximity is not supported for the ONTAP version. |

| Error Code | Description |
| --- | --- |
| 5376059 | Setting initiator proximity to a peer that is either the destination of an SVM DR relationship or in a Metrocluster configuration is not supported. |
| 26345672 | The specified SVM peering relationship was not found. |
| 26345673 | An SVM peering relationship between the initiator group's SVM and specified peer SVM was not found. |
| 26345675 | An SVM peering relationship UUID and name were specified and they do not refer to the same SVM peering relationship. |
| 26345680 | Supplied SVM peer is on the local cluster. The operation requires a peer on a remote cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interface

| Name | Type | Description |
|---|---|---|
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |

logins

| Name | Type | Description |
|---|---|---|
| connected | boolean | True if the initiator is currently logged in to this connection's interface. |
| interface | interface | |
| last_seen_time | string | The last time this initiator logged in. Logins not seen for 48 hours are cleared and not reported. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

connections

| Name | Type | Description |
| --- | --- | --- |
| logins | array[logins] | |
| node | node | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| connections | array[connections] | |

igroup

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator

in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

records

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

igroup_initiator

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br><br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group in which the initiator is found.<br><br>Note that this does not mean that the initiator cannot also be found in other initiator groups. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |
| records | array[records] | An array of initiators specified to add multiple initiators to an initiator group in a single API call. Not allowed when the `name` property is used. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an initiator group

DELETE `/protocols/san/igroups/{uuid}`

**Introduced In:** 9.6

Deletes an initiator group.

#### Related ONTAP commands

- `lun igroup delete`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the initiator group. |
| allow_delete_while_mapped | boolean | query | False | Allows the deletion of a mapped initiator group.<br><br>Deleting a mapped initiator group means that the LUNs, to which the initiator group is mapped, are no longer available to the initiators.<br>This might cause a disruption in the availability of data.<br><br>**This parameter should be used with caution.**<br><br>• Default value: |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254213 | The initiator group is mapped to one or more LUNs and `allow_delete_while_mapped` has not been specified. |
| 5374760 | An error was reported by the peer cluster while deleting a replicated initiator group. The specific error will be included as a nested error. |
| 5374852 | The initiator group does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an initiator group

GET `/protocols/san/igroups/{uuid}`

**Introduced In:** 9.6

Retrieves an initiator group.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `connectivity_tracking.*`
- `igroups.*`
- `lun_maps.*`
- `parent_igroups.*`

#### Related ONTAP commands

- `lun igroup show`
- `lun mapping show`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the initiator group. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroup_child] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[initiators] | The initiators that are members of the group or any group nested below this group. Optional in POST.<br><br>This property is mutually exclusive with the *igroups* property during POST.<br><br>During GET, this array contains initiators that are members of this group or any nested initiator groups below this group. When initiators of nested groups are returned, they include links to the initiator group that directly contains the initiator.<br><br>Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details. |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| replication | replication | Properties related to initiator group replication. |

| Name | Type | Description |
| --- | --- | --- |
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {
        "summary": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        }
      }
    ],
    "connection_state": "string",
    "required_nodes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  },
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "igroups": [
```

```
        null
      ],
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiators": [
    {
      "_links": {
        "connectivity_tracking": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "connectivity_tracking": {
        "connection_state": "string"
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "proximity": {
        "peer_svms": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          }
        ]
      }
    }
  ],
  "lun_maps": [
```

```
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "logical_unit_number": 0,
      "lun": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "lun1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "parent_igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "name": "igroup1",
      "parent_igroups": [
        null
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "portset": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "portset1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "protocol": "string",
  "replication": {
    "error": {
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "summary": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    },
    "peer_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "peer1",
      "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
    },
    "state": "string"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
```

```
    "firmware_revision": "9111",
    "product_id": "LUN C-Mode",
    "vendor_id": "NETAPP"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374852 | The initiator group does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state of the initiator group.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

| Name | Type | Description |
|---|---|---|
| summary | summary | A user friendly message describing the connection state of the initiator group. |

required_nodes

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

connectivity_tracking

An overview of the connections to ONTAP by the initiators in this group.

| Name | Type | Description |
|---|---|---|
| alerts | array[alerts] | |
| connection_state | string | Connection state. |
| required_nodes | array[required_nodes] | Nodes to which the initiators in this group should be connected to ensure reliable service. This is the collection of any node hosting a LUN mapped to this igroup as well as the HA partners of those nodes. |

igroup_child

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroup_child] | Further nested initiator groups. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

connectivity_tracking

A link to the initiator with connectivity information relevant to its membership of this initiator group.

| Name | Type | Description |
|---|---|---|
| href | string | |

self

A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group, the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group.

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| connectivity_tracking | connectivity_tracking | A link to the initiator with connectivity information relevant to its membership of this initiator group. |
| self | self | A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group, the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group. |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| connection_state | string | Connection state. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

initiators

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| name | string | The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.<br>An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters. |

| Name | Type | Description |
| --- | --- | --- |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

lun

The LUN to which the initiator group is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the LUN. |
| node | node | |
| uuid | string | The unique identifier of the LUN. |

lun_maps

A LUN map with which the initiator group is associated.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| logical_unit_number | integer | The logical unit number assigned to the LUN for initiators in the initiator group. |
| lun | lun | The LUN to which the initiator group is mapped. |

igroup_parent

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment available for use by the administrator. |
| name | string | The name of the initiator group. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| uuid | string | The unique identifier of the initiator group. |

portset

The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.

In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.

Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string ("").

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the portset. |
| uuid | string | The unique identifier of the portset. |

igroup

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | Indicates whether the reported igroup is on the local SVM or the peer SVM. When deleting a replicated igroup, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer igroup remains, the peer igroup is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

summary

A user friendly message describing the error.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
|---|---|---|
| igroup | igroup | |
| summary | summary | A user friendly message describing the error. |

peer_svm

The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to initiator group replication.

| Name | Type | Description |
|---|---|---|
| error | error | Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |

| Name | Type | Description |
| --- | --- | --- |
| peer_svm | peer_svm | The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`. |
| state | string | The state of the replication queue associated with this igroup. If this igroup is not in the replication queue, the state is reported as *ok*. If this igroup is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this igroup is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

Properties of the SCSI target to which the initiator group provides access.

| Name | Type | Description |
| --- | --- | --- |
| firmware_revision | string | The firmware revision of the SCSI target specific to the OS type of the initiator group. |

| Name | Type | Description |
|---|---|---|
| product_id | string | The product ID of the SCSI target. |
| vendor_id | string | The vendor ID of the SCSI target. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an initiator group

PATCH `/protocols/san/igroups/{uuid}`

**Introduced In:** 9.6

Updates an initiator group.

#### Related ONTAP commands

- `lun igroup modify`
- `lun igroup rename`
- `lun igroup bind`
- `lun igroup unbind`

#### Learn more

- DOC /protocols/san/igroups

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the initiator group. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| replication | replication | Properties related to initiator group replication. |
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

**Example request**

```
{
  "comment": "string",
  "connectivity_tracking": {
    "alerts": [
      {}
    ],
    "connection_state": "string",
    "required_nodes": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  },
  "lun_maps": [
    {
      "logical_unit_number": 0,
      "lun": {
        "name": "lun1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "parent_igroups": [
    {
      "comment": "string",
      "name": "igroup1",
      "parent_igroups": [
        null
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "portset": {
    "name": "portset1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "replication": {
```

```
    "error": {
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "peer_svm": {
      "name": "peer1",
      "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
    },
    "state": "string"
  },
  "target": {
    "firmware_revision": "9111",
    "product_id": "LUN C-Mode",
    "vendor_id": "NETAPP"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254264 | An attempt was made to bind a portset to an initiator group that is already bound to a portset. |
| 5373958 | An invalid initiator group name was supplied for a rename operation. |
| 5374023 | A rename operation failed because an initiator group with the same name already exists. |
| 5374027 | An attempt was made to bind a portset with no member network interfaces to the initiator group. |
| 5374028 | An attempt was made to bind a portset with an incompatible protocol to the initiator group. |

| Error Code | Description |
| --- | --- |
| 5374733 | An initiator is already in another initiator group with a conflicting operating system type. |
| 5374745 | An attempt was made to add an initiator group as a child to itself. |
| 5374746 | The cluster is currently running in a mixed version and nested initiator groups cannot be created until the effective cluster version reaches 9.9.1. |
| 5374747 | The cluster is currently running in a mixed version and initiator group comments cannot be created until the effective cluster version reaches 9.9.1. |
| 5374749 | An initiator group's replication peer SVM cannot be changed without first being cleared. |
| 5374759 | An error was reported by the peer cluster while modifying a replicated initiator group. The specific error will be included as a nested error. |
| 5374763 | An error was reported by the peer cluster while renaming a replicated initiator group. The specific error will be included as a nested error. |
| 5374765 | An initiator group cannot be replicated if it has unreplicated child initiator groups. |
| 5374766 | A replicated initiator group cannot be changed to unreplicated if it is the child of a replicated initiator group. |
| 5374770 | An unreplicated initiator group cannot be changed to replicated due to a conflict in its LUN maps. |
| 5374852 | The initiator group does not exist. |
| 5374868 | The initiator group was partially modified before an error was encountered while renaming the initiator group. |
| 5375258 | The igroup is already replicated to a different peer SVM. |
| 5376253 | Initiator group replication requires an effective cluster version of 9.15.1. |
| 5376255 | Initiator group replication requires the peer cluster to have an effective cluster version of 9.15.1. |
| 5376488 | An NVMe over Fabrics subsystem already exists with the requested name. |
| 6620376 | SVM peering information is unavailable. |
| 6620384 | The supplied SVMs are not peered. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

summary

A user friendly message describing the connection state of the initiator group.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

alerts

required_nodes

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

connectivity_tracking

An overview of the connections to ONTAP by the initiators in this group.

| Name | Type | Description |
|---|---|---|
| alerts | array[alerts] | |
| connection_state | string | Connection state. |

| Name | Type | Description |
| --- | --- | --- |
| required_nodes | array[required_nodes] | Nodes to which the initiators in this group should be connected to ensure reliable service. This is the collection of any node hosting a LUN mapped to this igroup as well as the HA partners of those nodes. |

igroup_child

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| igroups | array[igroup_child] | Further nested initiator groups. |

connectivity_tracking

A link to the initiator with connectivity information relevant to its membership of this initiator group.

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

self

A link to the initiator where mutations can be made. If the initiator is inherited from a nested initiator group, the link refers to the initiator in the nested initiator group. In this case, mutations of the initiator will be applied to all initiator groups referencing the same nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

connectivity_tracking

Overview of the initiator's connections to ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| connection_state | string | Connection state. |

igroup

The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

peer_svms

A reference to an SVM peer relationship.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

proximity

Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.

These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`.

| Name | Type | Description |
| --- | --- | --- |
| local_svm | boolean | A boolean that indicates if the initiator is proximal to the SVM of the containing initiator group. This is required for any POST or PATCH that includes the `proximity` sub-object. |
| peer_svms | array[peer_svms] | An array of remote peer SVMs to which the initiator is proximal. |

initiators

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | Overview of the initiator's connections to ONTAP.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| igroup | igroup | The initiator group that directly owns the initiator, which is where modification of the initiator is supported. This property will only be populated when the initiator is a member of a nested initiator group. |
| proximity | proximity | Properties that define to what SVMs the initiator is proximal. This information is used to properly report active optimized and active non-optimized network paths via ALUA. If no configuration has been specified for an initiator, the sub-object will not be present in GET.<br><br>These properties can be set via initiator group POST and PATCH and apply to all instances of the initiator in all initiator groups in the SVM and its peers. The `proximity` sub-object for an initiator is set in POST and PATCH in its entirety and replaces any previously set proximity for the initiator within the SVM for the initiator within the SVM. The `local_svm` property must always be set to `true` or `false` when setting the `proximity` property. To clear any previously set proximity, POST or PATCH the `proximity` object to `null`. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

lun

The LUN to which the initiator group is mapped.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. |
| node | node | |
| uuid | string | The unique identifier of the LUN. |

lun_maps

A LUN map with which the initiator group is associated.

| Name | Type | Description |
| --- | --- | --- |
| logical_unit_number | integer | The logical unit number assigned to the LUN for initiators in the initiator group. |
| lun | lun | The LUN to which the initiator group is mapped. |

igroup_parent

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | The name of the initiator group. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |
| uuid | string | The unique identifier of the initiator group. |

portset

The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.

In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.

Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string ("").

| Name | Type | Description |
|---|---|---|
| name | string | The name of the portset. |
| uuid | string | The unique identifier of the portset. |

igroup

| Name | Type | Description |
|---|---|---|
| local_svm | boolean | Indicates whether the reported igroup is on the local SVM or the peer SVM. When deleting a replicated igroup, the local copy is deleted first and then the peer copy is deleted. If the error is encountered between these two operations and only the peer igroup remains, the peer igroup is reported and the problem might need to be corrected on the peer cluster. |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

summary

A user friendly message describing the error.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

error

Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering

relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

| Name | Type | Description |
|---|---|---|
| igroup | igroup | |

peer_svm

The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

replication

Properties related to initiator group replication.

| Name | Type | Description |
| --- | --- | --- |
| error | error | Information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state. |
| peer_svm | peer_svm | The peered SVM to which the initiator group should be replicated. Optional in POST and PATCH. To clear any previously set replication peer, PATCH the `replication.peer_svm` object to `null`. |
| state | string | The state of the replication queue associated with this igroup. If this igroup is not in the replication queue, the state is reported as *ok*. If this igroup is in the replication queue, but no errors have been encountered, the state is reported as *replicating*. If this igroup is in the replication queue and the queue is blocked by an error, the state is reported as *error*. When in the *error* state, additional context is provided by the `replication.error` property. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

Properties of the SCSI target to which the initiator group provides access.

| Name | Type | Description |
|---|---|---|
| firmware_revision | string | The firmware revision of the SCSI target specific to the OS type of the initiator group. |
| product_id | string | The product ID of the SCSI target. |
| vendor_id | string | The vendor ID of the SCSI target. |

igroup

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPNs), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host endpoints.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the host initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator group may contain either initiators or other initiator groups, but not both simultaneously. When a parent initiator group is mapped, it inherits all of the initiators of any initiator groups nested below it. If any nested initiator group is modified to contain different initiators, the parent initiator groups inherit the change. A parent can have many nested initiator groups and an initiator group can be nested under multiple parents. Initiators can only be added or removed from the initiator group that directly contains them. The maximum supported depth of nesting is three layers.

Best practice when using nested initiator groups is to match host hierarchies. A single initiator group should correspond to a single host. If a LUN needs to be mapped to multiple hosts, the initiator groups representing those hosts should be aggregated into a parent initiator group and the LUN should be mapped to that initiator group. For multi-ported hosts, initiators have a comment property where the port corresponding to the initiator can be documented.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once. With the introduction of nestable initiator groups, best practice is to use the hierarchy such that an initiator is only a direct member of a single initiator group, and that initiator group can then be referenced by other initiator groups.

All initiators or nested initiator groups in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (*fcp*), iSCSI (*iscsi*), or both (*mixed*). Initiator groups within a nested hierarchy may not have conflicting protocols.

Zero or more initiators or nested initiator groups can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. Initiator groups containing other initiator groups report the aggregated list of initiators from all nested initiator groups, but modifications of the initiator list must be performed on the initiator group that directly contains the initiators. See DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name} for more details.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| connectivity_tracking | connectivity_tracking | An overview of the connections to ONTAP by the initiators in this group.<br>• readOnly: 1<br>• Introduced in: 9.11 |
| delete_on_unmap | boolean | An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to *false* when the initiator group is created. |

| Name | Type | Description |
| --- | --- | --- |
| lun_maps | array[lun_maps] | All LUN maps with which the initiator is associated.<br><br>If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| name | string | The name of the initiator group. Required in POST; optional in PATCH. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH. |
| parent_igroups | array[igroup_parent] | The initiator groups that contain this initiator group as a member. |

| Name | Type | Description |
| --- | --- | --- |
| portset | portset | The portset to which the initiator group is bound. Binding the initiator group to a portset restricts the initiators of the group to accessing mapped LUNs only through network interfaces in the portset.<br><br>In a nested initiator group hierarchy, only a portset bound to the initiator group at the same level at which it is mapped, applies; portsets bound to parent or child initiator groups are ignored.<br><br>Optional in POST and PATCH. To unbind a portset from the initiator group, PATCH the `portset` object to `null`, or PATCH `portset.name` to an empty string (""). |
| replication | replication | Properties related to initiator group replication. |
| supports_igroups | boolean | An initiator group may contain either initiators or other initiator groups, but not both simultaneously. This property is *true* when initiator groups can be added to this initiator group. The `initiators.name` property cannot be used to determine this via a query because it reports initiators inherited from nested igroups. |
| target | target | Properties of the SCSI target to which the initiator group provides access. |
| uuid | string | The unique identifier of the initiator group. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SAN initiators

### Manage SAN initiators

#### Overview

An initiator is a Fibre Channel (FC) world wide port name (WWPN), an iSCSI Qualified Name (IQN), or an iSCSI EUI (Extended Unique Identifier) that identifies a host endpoint. Initiators are collected into initiator groups (igroups) used to control which hosts can access specific LUNs. Initiators are also discovered as they log in to SAN network LIFs.

ONTAP supports configuration for an initiator. Configured properties apply to all uses of the initiator within an SVM. Although the same initiator may interact with multiple SVMs of a cluster, ONTAP treats initiator configuration as an SVM-scoped activity. For example, a comment may be set for an initiator in a specific SVM. The comment value applies to all uses of the initiator in the SVM including use in multiple initiator groups. But a different comment value may be specified for the same initiator in a different SVM.

The initiator REST API provides read-only access to properties of initiators.

An FC WWPN consists of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is *iqn.yyyy-mm.reverse_domain_name:any*. The iSCSI EUI format consists of the *eui.* prefix followed by 16 hexadecimal characters.

### Retrieve initiators

GET `/protocols/san/initiators`

**Introduced In:** 9.14

Retrieves initiators.

#### Related ONTAP commands

- `lun igroup initiator show`

#### Learn more

- DOC /protocols/san/initiators

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol | string | query | False | Filter by protocol |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| comment | string | query | False | Filter by comment<br>• maxLength: 254<br>• minLength: 0 |
| name | string | query | False | Filter by name |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[initiator] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "comment": "My initiator comment.",
      "name": "iqn.2018-02.com.netapp.iscsi:name1",
      "protocol": "iscsi",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

The SVM for which the initiator properties are configured.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

initiator

An initiator is a Fibre Channel (FC) world wide port name (WWPN), an iSCSI Qualified Name (IQN), or an iSCSI EUI (Extended Unique Identifier) that identifies a host endpoint. Initiators are collected into initiator groups (igroups) used to control which hosts can access specific LUNs. Initiators are also discovered as they log in to SAN network LIFs.

ONTAP supports configuration for an initiator. Configured properties apply to all uses of the initiator within an SVM. Although the same initiator may interact with multiple SVMs of a cluster, ONTAP treats initiator configuration as an SVM-scoped activity. For example, a comment may be set for an initiator in a specific SVM. The comment value applies to all uses of the initiator in the SVM including use in multiple initiator groups. But a different comment value may be specified for the same initiator in a different SVM.

| Name | Type | Description |
|---|---|---|
| comment | string | A user-specified comment. |
| name | string | The name of the initiator. |
| protocol | string | The protocol of the initiator. |
| svm | svm | The SVM for which the initiator properties are configured. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an initiator

GET `/protocols/san/initiators/{svm.uuid}/{name}`

**Introduced In:** 9.14

Retrieves an initiator.

#### Related ONTAP commands

- `lun igroup initiator show`

#### Learn more

- DOC /protocols/san/initiators

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which the initiator properties are configured. |
| name | string | path | True | The name of the initiator. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| comment | string | A user-specified comment. |
| name | string | The name of the initiator. |
| protocol | string | The protocol of the initiator. |
| svm | svm | The SVM for which the initiator properties are configured. |

### Example response

```
{
  "comment": "My initiator comment.",
  "name": "iqn.2018-02.com.netapp.iscsi:name1",
  "protocol": "iscsi",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

The SVM for which the initiator properties are configured.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage iSCSI credentials

## Manage iSCSI credentials

### Overview

An iSCSI credentials object defines authentication credentials to be used between an initiator and ONTAP. It identifies an authentication type, user names, and passwords that must be used to authenticate a specific initiator.

The iSCSI credentials REST API allows you to create, update, delete, and discover iSCSI credentials.

### How iSCSI authentication works

An iSCSI credentials object defines the authentication credentials to be used between an initiator and ONTAP. While establishing an iSCSI connection, the initiator sends a login request to ONTAP to begin an iSCSI session. ONTAP then either permits or denies the login request, or determines that a login is not required.

For an initiator, you can specify an authentication type, user names and passwords, and a whitelist of optional network addresses from which the initiator is allowed to connect.

### iSCSI authentication methods

- Challenge-Handshake Authentication Protocol (CHAP) - The initiator logs in using a CHAP user name and password. There are two types of CHAP user names and passwords:
  - Inbound - ONTAP authenticates the initiator. Inbound settings are required if you are using CHAP authentication.
  - Outbound - These are optional credentials to enable the initiator to authenticate ONTAP. You can use credentials only if inbound credentials are also being used.
- deny - The initiator is denied access to ONTAP.
- none - ONTAP does not require authentication for the initiator.
  The CHAP inbound/outbound password can be any valid string or an even number of valid hexadecimal digits preceded by '0X' or '0x'.

### Initiator address list

The initiator address list is a way to specify valid IP addresses from which the initiator is allowed to connect. If the list is specified and the source address of an iSCSI connection is not in the list, the connection is rejected. Initiator addresses can be specified in either IPv4 or IPv6 format and in one of two forms:

- Range

```
{
  "start": "192.168.0.0",
  "end": "192.168.0.255"
}
```

- Mask

```
{
  "address": "192.168.0.0",
  "netmask": "24"
}
```

#### Initiator "default"

The default iSCSI authentication definition is created when the iSCSI service is created. An iSCSI credentials object with *default* as the initiator name identifies the default authentication for an SVM. The default credentials are used for any initiator that does not have specific iSCSI credentials. The default iSCSI authentication method is *none*, but can be changed to *deny* or *CHAP*. The default credentials object does not support an initiator address list.

#### Examples

### Creating iSCSI credentials requiring no authentication

```
# The API:
POST /api/protocols/san/iscsi/credentials

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/iscsi/credentials' -H
'Accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"initiator": "iqn.1992-08.com.netapp:initiator1", "authentication_type":
"none" }'
```

---

### Creating iSCSI credentials using CHAP inbound authentication

```
# The API:
POST /api/protocols/san/iscsi/credentials

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/iscsi/credentials' -H
'Accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"initiator": "iqn.1992-08.com.netapp:initiator2", "authentication_type":
"CHAP", "chap": { "inbound": { "user": "user1", "password": "password1" }
} }'
```

---

### Retrieving all properties of all iSCSI credentials

The `fields` query parameter is used to request all iSCSI credentials properties.

Passwords are not included in the GET output.

```
# The API:
GET /api/protocols/san/iscsi/credentials

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials?fields=*' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "default",
    "authentication_type": "none",
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-
8370-005056b48fd2/default"
      }
    }
  },
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "iqn.1992-08.com.netapp:initiator1",
    "authentication_type": "none",
    "_links": {
      "self": {
```

```
        "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-
8370-005056b48fd2/iqn.1992-08.com.netapp:initiator1"
      }
    }
  },
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "iqn.1992-08.com.netapp:initiator2",
    "authentication_type": "chap",
    "chap": {
      "inbound": {
        "user": "user1"
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-
8370-005056b48fd2/iqn.1992-08.com.netapp:initiator2"
      }
    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "default",
    "authentication_type": "none",
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-
8370-005056b48fd2/default"
      }
```

```
    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "iqn.1992-08.com.netapp:initiator2",
    "authentication_type": "none",
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-
8370-005056b48fd2/iqn.1992-08.com.netapp:initiator2"
      }
    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "iqn.1992-08.com.netapp:initiator3",
    "authentication_type": "deny",
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-
8370-005056b48fd2/iqn.1992-08.com.netapp:initiator3"
      }
    }
  }
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/credentials?fields=*"
  }
```

```
}
}
```

### Retrieving specific iSCSI credentials

```
# The API:
GET /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'Accept:
application/hal+json'

# The response:
{
"svm": {
  "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
  "name": "svm2",
  "_links": {
    "self": {
      "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
    }
  }
},
"initiator": "iqn.1992-08.com.netapp:initiator2",
"authentication_type": "chap",
"chap": {
  "inbound": {
    "user": "user1"
  }
},
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2"
  }
}
}
```

#### Updating the authentication type of iSCSI credentials

```
# The API:
PATCH /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

# The call:
curl -X PATCH 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'Accept:
application/hal+json' -d '{ "authentication_type": "chap", "chap": {
"inbound": { "user": "user1", "password": "password1" } } }'
```

---

#### Updating the initiator address list of iSCSI credentials

```
# The API:
PATCH /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

# The call:
curl -X PATCH 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'Accept:
application/hal+json' -d '{ "initiator_address": { "ranges": [ { "start":
"192.168.0.0", "end": "192.168.255.255" } ] } }'
```

---

#### Deleting iSCSI credentials

```
# The API:
DELETE /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

# The call:
curl -X DELETE 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'Accept:
application/hal+json'
```

#### Retrieve iSCSI credentials

GET `/protocols/san/iscsi/credentials`

**Introduced In:** 9.6

Retrieves iSCSI credentials.

#### Related ONTAP commands

- `vserver iscsi security show`

#### Learn more

- DOC /protocols/san/iscsi/credentials

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| chap.outbound.user | string | query | False | Filter by chap.outbound.user<br>• maxLength: 128<br>• minLength: 1 |
| chap.inbound.user | string | query | False | Filter by chap.inbound.user<br>• maxLength: 128<br>• minLength: 1 |
| authentication_type | string | query | False | Filter by authentication_type |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| initiator_address.ranges.family | string | query | False | Filter by initiator_address.ranges.family |
| initiator_address.ranges.start | string | query | False | Filter by initiator_address.ranges.start |
| initiator_address.ranges.end | string | query | False | Filter by initiator_address.ranges.end |
| initiator_address.masks.family | string | query | False | Filter by initiator_address.masks.family |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| initiator_address.masks.address | string | query | False | Filter by initiator_address.masks.address |
| initiator_address.masks.netmask | string | query | False | Filter by initiator_address.masks.netmask |
| initiator | string | query | False | Filter by initiator |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[iscsi_credentials] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_type": "string",
      "chap": {
        "inbound": {
          "user": "string"
        },
        "outbound": {
          "user": "string"
        }
      },
      "initiator": "iqn.1998-01.com.corp.iscsi:name1",
      "initiator_address": {
        "masks": [
          {
            "address": "10.10.10.7",
            "family": "string",
            "netmask": "24"
          }
        ],
        "ranges": [
          {
            "end": "10.10.10.7",
            "family": "string",
            "start": "10.10.10.7"
          }
        ]
      },
      "svm": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

inbound

Inbound CHAP credentials.

| Name | Type | Description |
|---|---|---|
| user | string | The inbound CHAP user name. Optional in POST and PATCH. |

outbound

Output CHAP credentials.

To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH.

| Name | Type | Description |
|---|---|---|
| user | string | The outbound CHAP user name. Optional in POST and PATCH.<br><br>To clear previously set outbound CHAP credentials, set this property to an empty string in PATCH. |

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

| Name | Type | Description |
| --- | --- | --- |
| inbound | inbound | Inbound CHAP credentials. |
| outbound | outbound | Output CHAP credentials.<br><br>To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_address_range

IP address range

| Name | Type | Description |
| --- | --- | --- |
| end | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| start | string | IPv4 or IPv6 address |

initiator_address

Initiator address ranges.

| Name | Type | Description |
| --- | --- | --- |
| masks | array[ip_info] | |
| ranges | array[ip_address_range] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

iscsi_credentials

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator | string | The iSCSI initiator to which the credentials apply. Required in POST. |
| initiator_address | initiator_address | Initiator address ranges. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create iSCSI credentials

POST `/protocols/san/iscsi/credentials`

**Introduced In:** 9.6

Creates iSCSI credentials.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the iSCSI credentials.
- `initiator` - Initiator for which the iSCSI credentials are to be created.
- `authentication_type` - Type of authentication to use for the credentials.

#### Recommended optional properties

- `chap.inbound.user` - In-bound CHAP authentication user name.
- `chap.inbound.password` - In-bound CHAP authentication password.
- `chap.outbound.user` - Out-bound CHAP authentication user name.
- `chap.outbound.password` - Out-bound CHAP authentication password.

#### Related ONTAP commands

- `vserver iscsi security create`

#### Learn more

- DOC /protocols/san/iscsi/credentials

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator | string | The iSCSI initiator to which the credentials apply. Required in POST. |
| initiator_address | initiator_address | Initiator address ranges. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "password": "string",
      "user": "string"
    },
    "outbound": {
      "password": "string",
      "user": "string"
    }
  },
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": [
      {
        "address": "10.10.10.7",
        "netmask": "24"
      }
    ],
    "ranges": [
      {
        "end": "10.10.10.7",
        "start": "10.10.10.7"
      }
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |

| Name | Type | Description |
| --- | --- | --- |
| records | array[iscsi_credentials] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "authentication_type": "string",
      "chap": {
        "inbound": {
          "password": "string",
          "user": "string"
        },
        "outbound": {
          "password": "string",
          "user": "string"
        }
      },
      "initiator": "iqn.1998-01.com.corp.iscsi:name1",
      "initiator_address": {
        "masks": [
          {
            "address": "10.10.10.7",
            "netmask": "24"
          }
        ],
        "ranges": [
          {
            "end": "10.10.10.7",
            "start": "10.10.10.7"
          }
        ]
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 2621706 | Both the SVM UUID and SVM name were supplied, but they do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373969 | A non-empty qualifier is required after the prefix. An example of a valid IQN is *iqn.1995-08.com.example:string*. |
| 5373970 | The IQN prefix is invalid. The correct IQN prefix is *iqn*. An example of a valid IQN is *iqn.1995-08.com.example:string*. |
| 5373971 | The date field is invalid. A valid date field is *yyyy-mm*. An example of a valid IQN is *iqn.1995-08.com.example:string*. |
| 5373972 | The naming authority and string fields can contain only the characters *a-z*, *0-9*, ., :, and -. |
| 5373977 | The EUI-64 identifier field must be exactly 16 hexadecimal digits. |
| 5373978 | The EUI formatted initiator name supplied is invalid. A valid EUI format is *eui.XXXXXXXXXXXXXXXX*, where X is a hexadecimal digit. |
| 5373997 | The initiator name supplied in invalid. The valid initiator name formats are *iqn.1995-08.com.example:string* or *eui.0123456789abcdef*. |
| 5374078 | The iSCSI service does not exist. |
| 5374142 | An iSCSI security credential already exists for the specified initiator. |
| 5374145 | The iSCSI security password must contain an even number of valid hex digits. |

| Error Code | Description |
| --- | --- |
| 5374147 | The CHAP inbound and outbound passwords must be different. |
| 5374149 | The inbound user and password properties are required for CHAP authentication. |
| 5374150 | Outbound CHAP authentication requires an outbound password. |
| 5374855 | The value for property `initiator_address.ranges.start` is greater than the value for property `initiator_address.ranges.end`. |
| 5374856 | The value for property `initiator_address.ranges.start` does not belong to the same IP address family as the value for property `initiator_address.ranges.end`. |
| 5374900 | Setting the CHAP authentication properties are not supported with authentication types *none* or *deny*. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

inbound

Inbound CHAP credentials.

| Name | Type | Description |
|---|---|---|
| password | string | The inbound CHAP password. Write-only; optional in POST and PATCH. |
| user | string | The inbound CHAP user name. Optional in POST and PATCH. |

outbound

Output CHAP credentials.

To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH.

| Name | Type | Description |
|---|---|---|
| password | string | The outbound CHAP password. Write-only; optional in POST and PATCH. |
| user | string | The outbound CHAP user name. Optional in POST and PATCH.<br><br>To clear previously set outbound CHAP credentials, set this property to an empty string in PATCH. |

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

| Name | Type | Description |
|---|---|---|
| inbound | inbound | Inbound CHAP credentials. |

| Name | Type | Description |
| --- | --- | --- |
| outbound | outbound | Output CHAP credentials.<br><br>To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_address_range

IP address range

| Name | Type | Description |
| --- | --- | --- |
| end | string | IPv4 or IPv6 address |
| start | string | IPv4 or IPv6 address |

initiator_address

Initiator address ranges.

| Name | Type | Description |
| --- | --- | --- |
| masks | array[ip_info] | |
| ranges | array[ip_address_range] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

iscsi_credentials

| Name | Type | Description |
| --- | --- | --- |
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator | string | The iSCSI initiator to which the credentials apply. Required in POST. |
| initiator_address | initiator_address | Initiator address ranges. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete iSCSI credentials

DELETE `/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`

**Introduced In:** 9.6

Deletes specified iSCSI credentials.

#### Related ONTAP commands

- `vserver iscsi security delete`

#### Learn more

- DOC /protocols/san/iscsi/credentials

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| svm.uuid | string | path | True | The unique identifier of an SVM. |
| initiator | string | path | True | The iSCSI initiator of the credentials object. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2621462 | An SVM with the specified UUID does not exist. |
| 2621706 | Both the SVM UUID and SVM name were supplied, but they do not refer to the same SVM. |

| Error Code | Description |
| --- | --- |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374148 | The default security credential cannot be deleted for an SVM. |
| 5374895 | The iSCSI security credential does not exist on the specified SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve specific iSCSI credentials

GET `/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`

**Introduced In:** 9.6

Retrieves specified iSCSI credentials.

#### Related ONTAP commands

- `vserver iscsi security show`

#### Learn more

- DOC /protocols/san/iscsi/credentials

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of an SVM. |
| initiator | string | path | True | The iSCSI initiator of the credentials object. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator | string | The iSCSI initiator to which the credentials apply. Required in POST. |
| initiator_address | initiator_address | Initiator address ranges. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "user": "string"
    },
    "outbound": {
      "user": "string"
    }
  },
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": [
      {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      }
    ],
    "ranges": [
      {
        "end": "10.10.10.7",
        "family": "string",
        "start": "10.10.10.7"
      }
    ]
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

inbound

Inbound CHAP credentials.

| Name | Type | Description |
|---|---|---|
| user | string | The inbound CHAP user name. Optional in POST and PATCH. |

outbound

Output CHAP credentials.

To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH.

| Name | Type | Description |
|---|---|---|
| user | string | The outbound CHAP user name. Optional in POST and PATCH.<br><br>To clear previously set outbound CHAP credentials, set this property to an empty string in PATCH. |

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

| Name | Type | Description |
|---|---|---|
| inbound | inbound | Inbound CHAP credentials. |

| Name | Type | Description |
| --- | --- | --- |
| outbound | outbound | Output CHAP credentials.<br><br>To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_address_range

IP address range

| Name | Type | Description |
| --- | --- | --- |
| end | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| start | string | IPv4 or IPv6 address |

initiator_address

Initiator address ranges.

| Name | Type | Description |
| --- | --- | --- |
| masks | array[ip_info] | |
| ranges | array[ip_address_range] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update iSCSI credentials

PATCH `/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`

**Introduced In:** 9.6

Updates specified iSCSI credentials.

#### Related ONTAP commands

- `vserver iscsi security add-initiator-address-ranges`
- `vserver iscsi security default`
- `vserver iscsi security modify`
- `vserver iscsi security remove-initiator-address-ranges`

#### Learn more

- DOC /protocols/san/iscsi/credentials

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of an SVM. |
| initiator | string | path | True | The iSCSI initiator of the credentials object. |
| add_initiator_addresses | boolean | query | False | If *true*, the initiator addresses in the body merge into the existing addresses in the iSCSI security object rather than replace the existing addresses.<br>• Default value: |
| remove_initiator_addresses | boolean | query | False | If *true*, the initiator addresses in the body are removed from the existing addresses in the iSCSI security object rather than replace the existing addresses.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator_address | initiator_address | Initiator address ranges. |

#### Example request

```
{
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "password": "string",
      "user": "string"
    },
    "outbound": {
      "password": "string",
      "user": "string"
    }
  },
  "initiator_address": {
    "masks": [
      {
        "address": "10.10.10.7",
        "netmask": "24"
      }
    ],
    "ranges": [
      {
        "end": "10.10.10.7",
        "start": "10.10.10.7"
      }
    ]
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |

| Error Code | Description |
| --- | --- |
| 2621706 | Both the SVM UUID and SVM name were supplied, but they do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374145 | The iSCSI security password must contain an even number of valid hex digits. |
| 5374147 | The CHAP inbound and outbound passwords must be different. |
| 5374149 | The inbound user and password properties are required for CHAP authentication. |
| 5374150 | Outbound CHAP authentication requires an outbound password. |
| 5374155 | The functionality is not supported for the default security credential. |
| 5374855 | The value for property `initiator_address.ranges.start` is greater than the value for property `initiator_address.ranges.end`. |
| 5374856 | The value for property `initiator_address.ranges.start` does not belong to the same IP address family as the value for property `initiator_address.ranges.end`. |
| 5374895 | The iSCSI security credential does not exist on the specified SVM. |
| 5374900 | Setting the CHAP authentication properties are not supported with authentication types *none* or *deny*. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

inbound

Inbound CHAP credentials.

| Name | Type | Description |
|---|---|---|
| password | string | The inbound CHAP password. Write-only; optional in POST and PATCH. |
| user | string | The inbound CHAP user name. Optional in POST and PATCH. |

outbound

Output CHAP credentials.

To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH.

| Name | Type | Description |
|---|---|---|
| password | string | The outbound CHAP password. Write-only; optional in POST and PATCH. |
| user | string | The outbound CHAP user name. Optional in POST and PATCH.<br><br>To clear previously set outbound CHAP credentials, set this property to an empty string in PATCH. |

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

| Name | Type | Description |
|---|---|---|
| inbound | inbound | Inbound CHAP credentials. |

| Name | Type | Description |
| --- | --- | --- |
| outbound | outbound | Output CHAP credentials.<br><br>To clear previously set outbound CHAP credentials, set property `chap.outbound.user` to an empty string in PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

ip_address_range

IP address range

| Name | Type | Description |
| --- | --- | --- |
| end | string | IPv4 or IPv6 address |
| start | string | IPv4 or IPv6 address |

initiator_address

Initiator address ranges.

| Name | Type | Description |
| --- | --- | --- |
| masks | array[ip_info] | |
| ranges | array[ip_address_range] | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

iscsi_credentials

| Name | Type | Description |
| --- | --- | --- |
| authentication_type | string | The iSCSI authentication type. Required in POST; optional in PATCH. |
| chap | chap | Challenge-Handshake Authentication Protocol (CHAP) credentials. |
| initiator_address | initiator_address | Initiator address ranges. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Manage iSCSI services

## Manage iSCSI services

### Overview

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

The iSCSI service REST API allows you to create, update, delete, and discover iSCSI services for SVMs.

### Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

### Examples

## Creating an iSCSI service for an SVM

The simplest way to create an iSCSI service is to specify only the SVM, either by name or UUID. By default, the new iSCSI service is enabled and uses the SVM name as its target alias.

In this example, the `return_records` query parameter is used to retrieve the new iSCSI service object in the REST response.

```
# The API:
POST /api/protocols/san/iscsi/services

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/iscsi/services?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" } }'

# The response:
{
"num_records": 1,
"records": [
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "enabled": true,
    "target": {
      "name": "iqn.1992-
08.com.netapp:sn.19d04b8e94d711e88370005056b48fd2:vs.4",
      "alias": "svm1"
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-
8370-005056b48fd2"
      }
    }
  }
]
}
```

#### Retrieving the iSCSI services for all SVMs in the cluster

```
# The API:
GET /api/protocols/san/iscsi/services
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-
8370-005056b48fd2"
      }
    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/services/25f617cf-94d7-11e8-
8370-005056b48fd2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/services"
```

```
  }
}
}
```

#### Retrieving details for a specific iSCSI service

The iSCSI service is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/san/iscsi/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services/19d04b8e-
94d7-11e8-8370-005056b48fd2' -H 'Accept: application/hal+json'

# The response:
{
"svm": {
  "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
    }
  }
},
"enabled": true,
"target": {
  "name": "iqn.1992-
08.com.netapp:sn.19d04b8e94d711e88370005056b48fd2:vs.4",
  "alias": "svm1"
},
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-
005056b48fd2"
  }
}
}
```

### Disabling an iSCSI service

Disabling an iSCSI service shuts down all active iSCSI sessions for the SVM and prevents the creation of new iSCSI sessions.

The iSCSI service to update is identified by the UUID of its SVM.

```
# The API:
PATCH /api/protocols/san/iscsi/services/{svm.uuid}

# The call:
curl -X PATCH 'https://<mgmt-
ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2'
-H 'Accept: application/hal+json' -d '{ "enabled": "false" }'
```

You can retrieve the iSCSI service to confirm the change.

In this example, the `fields` query parameter is used to limit the response to the `enabled` property and iSCSI service identifiers.

```
# The API:
GET /api/protocols/san/iscsi/services/{svm.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services/19d04b8e-
94d7-11e8-8370-005056b48fd2?fields=enabled' -H 'Accept:
application/hal+json'

# The response:
{
"svm": {
  "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
    }
  }
},
"enabled": false,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-
005056b48fd2"
  }
}
}
```

#### Deleting an iSCSI service

The iSCSI service must be disabled before it can be deleted.

The iSCSI service to be deleted is identified by the UUID of its SVM.

```
# The API:
DELETE /api/protocols/san/iscsi/services/{svm.uuid}

# The call:
curl -X DELETE 'https://<mgmt-
ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2'
-H 'Accept: application/hal+json'
```

### Retrieve iSCSI services

GET `/protocols/san/iscsi/services`

**Introduced In:** 9.6

Retrieves iSCSI services.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `statistics.*`
- `metric.*`

#### Related ONTAP commands

- `vserver iscsi show`

#### Learn more

- DOC /protocols/san/iscsi/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br>• Introduced in: 9.7 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br><br>• Introduced in: 9.7 |
| statistics.status | string | query | False | Filter by statistics.status<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br><br>• Introduced in: 9.7 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br><br>• Introduced in: 9.7 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.7 |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.7 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.7 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.7 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.7 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.7 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.7 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br>• Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status<br>• Introduced in: 9.7 |
| metric.duration | string | query | False | Filter by metric.duration<br>• Introduced in: 9.7 |
| target.name | string | query | False | Filter by target.name<br>• maxLength: 128<br>• minLength: 1 |
| target.alias | string | query | False | Filter by target.alias<br>• maxLength: 128<br>• minLength: 1 |
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[iscsi_service] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "statistics": {
        "iops_raw": {
```

```
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "alias": "svm1",
        "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|---|---|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |

| Name | Type | Description |
| --- | --- | --- |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
| --- | --- | --- |
| alias | string | The iSCSI target alias of the iSCSI service.<br><br>The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.<br><br>Optional in POST and PATCH. In POST, this defaults to the name of the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iscsi_service

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an iSCSI service

POST `/protocols/san/iscsi/services`

**Introduced In:** 9.6

Creates an iSCSI service.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the iSCSI service.

#### Related ONTAP commands

- `vserver iscsi create`

#### Learn more

- DOC /protocols/san/iscsi/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
  }
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[iscsi_service] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "alias": "svm1",
        "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1115127 | The cluster lacks a valid iSCSI license. |
| 2621462 | The supplied SVM does not exist. |
| 2621507 | The iSCSI protocol is not allowed for the specified SVM. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373966 | An iSCSI service cannot be created in an SVM that is configured for NVMe. |

| Error Code | Description |
| --- | --- |
| 5374077 | An iSCSI service already exists for the specified SVM. |
| 5374893 | The SVM is stopped. The SVM must be running to create an iSCSI service. |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| alias | string | The iSCSI target alias of the iSCSI service.<br><br>The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.<br><br>Optional in POST and PATCH. In POST, this defaults to the name of the SVM. |
| name | string | The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iscsi_service

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an iSCSI service**

DELETE `/protocols/san/iscsi/services/{svm.uuid}`

**Introduced In:** 9.6

Deletes an iSCSI service. An iSCSI service must be disabled before it can be deleted.

#### Related ONTAP commands

- `vserver iscsi delete`

#### Learn more

- DOC /protocols/san/iscsi/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which to delete the iSCSI service. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 5373960 | The iSCSI service is enabled. The iSCSI service must be disabled before it can be deleted. |
| 5374078 | The SVM does not have an iSCSI service. |
| 5376452 | Service POST and DELETE are not supported on ASA r2. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an iSCSI service

GET `/protocols/san/iscsi/services/{svm.uuid}`

**Introduced In:** 9.6

Retrieves an iSCSI service.

#### Related ONTAP commands

- `vserver iscsi show`

#### Learn more

- DOC /protocols/san/iscsi/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which to retrieve the iSCSI service. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| metric | metric | Performance numbers, such as IOPS latency and throughput, for SVM protocols. |

| Name | Type | Description |
| --- | --- | --- |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target | target | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
```

```
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 5374078 | The SVM does not have an iSCSI service. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| alias | string | The iSCSI target alias of the iSCSI service.<br><br>The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.<br><br>Optional in POST and PATCH. In POST, this defaults to the name of the SVM. |
| name | string | The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an iSCSI service

PATCH `/protocols/san/iscsi/services/{svm.uuid}`

**Introduced In:** 9.6

Updates an iSCSI service.

#### Related ONTAP commands

- `vserver iscsi modify`
- `vserver iscsi start`
- `vserver iscsi stop`

#### Learn more

- DOC /protocols/san/iscsi/services

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM for which to update the iSCSI service. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| target | target | |

#### Example request

```
{
  "target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |
| 5374078 | The SVM does not have an iSCSI service. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

target

| Name | Type | Description |
|---|---|---|
| alias | string | The iSCSI target alias of the iSCSI service.<br><br>The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.<br><br>Optional in POST and PATCH. In POST, this defaults to the name of the SVM. |
| name | string | The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.<br><br>If required, the target name can be modified using the ONTAP command line.<br><br>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2<br>• maxLength: 128<br>• minLength: 1<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iscsi_service

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.<br><br>Optional in POST and PATCH. The default setting is *true* (enabled) in POST. |
| target | target | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for the iSCSI protocol of an SVM**

GET `/protocols/san/iscsi/services/{svm.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for the iSCSI protocol service of an SVM.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| status | string | query | False | Filter by status |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| timestamp | string | query | False | Filter by timestamp |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.total | integer | query | False | Filter by throughput.total |
| duration | string | query | False | Filter by duration |
| svm.uuid | string | path | True | The unique identifier of the SVM. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "svm": {
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for the iSCSI protocol service of an SVM for a specific time**

GET `/protocols/san/iscsi/services/{svm.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical performance metrics for the iSCSI protocol service of an SVM for a specific time.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM. |
| timestamp | string | path | True | The timestamp of the performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| svm | svm | |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "svm": {
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## View iSCSI sessions

### View iSCSI sessions

#### Overview

An iSCSI session is one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

The iSCSI sessions REST API provides information about iSCSI initiators that have successfully logged in to ONTAP.

#### Examples

### Retrieving all iSCSI sessions

```
# The API:
GET /api/protocols/san/iscsi/sessions

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/san/iscsi/sessions" -H
"Accept: application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "target_portal_group": "iscsi_lif1",
    "tsih": 10,
    "_links": {
      "self": {
```

```
        "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-
7575-ada21efcaf16/iscsi_lif1/10"
      }
    }
  },
  {
    "svm": {
      "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "target_portal_group": "iscsi_lif2",
    "tsih": 11,
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/sessions/b009a9e7-4081-b576-
7575-ada21efcaf16/iscsi_lif2/11"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/sessions"
  }
}
}
```


---

**Retrieving all of the iSCSI sessions under the target portal group *iscsi_lif1***

The `tpgroup` query parameter is used to perform the query.

```
# The API:
GET /api/protocols/san/iscsi/sessions

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/san/iscsi/sessions?tpgroup=iscsi_lif1" -H "Accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
        }
      }
    },
    "target_portal_group": "iscsi_lif1",
    "tsih": 10,
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-
7575-ada21efcaf16/iscsi_lif1/10"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/sessions"
  }
}
}
```

#### Retrieving an iSCSI session

```
# The API:
GET
```

```
/api/protocols/san/iscsi/sessions/{svm.uuid}/{target_portal_group}/{tsih}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/san/iscsi/sessions/a009a9e7-
4081-b576-7575-ada21efcaf16/iscsi_lif1/10" -H "Accept:
application/hal+json"

# The response:
{
"svm": {
  "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
    }
  }
},
"target_portal_group": "iscsi_lif1",
"tsih": 10,
"initiator": {
  "name": "iqn.1994-05.com.example:string",
  "comment": "Example information about this initiator"
},
"isid": "61:62:63:64:65:00",
"target_portal_group_tag": 1027,
"connections": [
  {
    "cid": 1,
    "authentication_type": "chap",
    "initiator_address": {
      "address": "10.224.123.85",
      "port": 43827
    },
    "interface": {
      "name": "iscsi_lif1",
      "uuid": "c15439b4-dbb4-11e8-90ac-005056bba882",
      "ip": {
        "address": "192.168.0.1",
        "port": 3260
      },
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/c15439b4-dbb4-11e8-90ac-
005056bba882"
        }
```

```
      }
    }
  }
],
"igroups": [
  {
    "uuid": "af7838cd-f993-4faf-90b7-5524787ae1e8",
    "name": "igroup1",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/af7838cd-f993-4faf-90b7-
5524787ae1e8"
      }
    }
  },
  {
    "uuid": "bf7838cd-f993-4faf-90b7-5524787ae1e8",
    "name": "igroup2",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/bf7838cd-f993-4faf-90b7-
5524787ae1e8"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-7575-
ada21efcaf16/iscsi_lif1/10"
  }
}
}
```


### Retrieve iSCSI sessions

GET `/protocols/san/iscsi/sessions`

**Introduced In:** 9.6

Retrieves iSCSI sessions.

#### Related ONTAP commands

- `vserver iscsi connection show`
- `vserver iscsi session parameter show`

- `vserver iscsi session show`

**Learn more**

- DOC /protocols/san/iscsi/sessions

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| tsih | integer | query | False | Filter by tsih |
| connections.authentication_type | string | query | False | Filter by connections.authentication_type |
| connections.initiator_address.port | integer | query | False | Filter by connections.initiator_address.port |
| connections.initiator_address.address | string | query | False | Filter by connections.initiator_address.address |
| connections.interface.ip.port | integer | query | False | Filter by connections.interface.ip.port<br>• Max value: 65536<br>• Min value: 1 |
| connections.interface.ip.address | string | query | False | Filter by connections.interface.ip.address |
| connections.interface.name | string | query | False | Filter by connections.interface.name |
| connections.interface.uuid | string | query | False | Filter by connections.interface.uuid |
| connections.cid | integer | query | False | Filter by connections.cid |
| isid | string | query | False | Filter by isid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| target_portal_group | string | query | False | Filter by target_portal_group |
| igroups.name | string | query | False | Filter by igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| igroups.uuid | string | query | False | Filter by igroups.uuid |
| target_portal_group_tag | integer | query | False | Filter by target_portal_group_tag |
| initiator.comment | string | query | False | Filter by initiator.comment<br>• Introduced in: 9.9 |
| initiator.name | string | query | False | Filter by initiator.name |
| initiator.alias | string | query | False | Filter by initiator.alias |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[iscsi_session] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "connections": [
        {
          "_links": {
            "next": {
              "href": "/api/resourcelink"
            },
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "authentication_type": "string",
          "cid": 0,
          "initiator_address": {
            "address": "10.10.10.7",
            "port": 55432
          },
          "interface": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "ip": {
              "address": "10.10.10.7",
              "port": 3260
            },
            "name": "lif1",
```

```
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "igroups": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "initiator": {
        "alias": "initiator_alias1",
        "comment": "This is an iSCSI initiator for host 5",
        "name": "iqn.1992-01.example.com:string"
      },
      "isid": "61:62:63:64:65:00",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target_portal_group": "tpgroup1",
      "target_portal_group_tag": 0,
      "tsih": 0
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

initiator_address

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

| Name | Type | Description |
|---|---|---|
| address | string | The TCP IPv4 or IPv6 address of the initiator end of the iSCSI connection. |
| port | integer | The TCP port number of the initiator end of the iSCSI connection. |

ip

The IP information. ONTAP only supports port 3260.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| port | integer | The TCP port number of the iSCSI access endpoint. |

interface

The network interface information for the target end of the connection.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | The IP information. ONTAP only supports port 3260.<br>• readOnly: 1<br>• Introduced in: 9.6 |
| name | string | The name of the interface. |
| uuid | string | The UUID that uniquely identifies the interface. |

iscsi_connection

An active iSCSI connection.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication_type | string | The iSCSI authentication type used to establish the connection. |
| cid | integer | The identifier of the connection within the session. |
| initiator_address | initiator_address | The TCP socket information for the initiator end of the connection. This is useful for network packet debugging. |
| interface | interface | The network interface information for the target end of the connection. |

igroups

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiator

The initiator that created the session.

| Name | Type | Description |
|---|---|---|
| alias | string | The initiator alias. |
| comment | string | A comment available for use by the administrator. This is modifiable from the initiator REST endpoint directly. See PATCH /protocols/san/igroups/{igroup.uuid}/initiators/{name} . |
| name | string | The world wide unique name of the initiator. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

iscsi_session

An iSCSI session is one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| connections | array[iscsi_connection] | The iSCSI connections that make up the iSCSI session. |
| igroups | array[igroups] | The initiator groups in which the initiator is a member. |
| initiator | initiator | The initiator that created the session. |

| Name | Type | Description |
|---|---|---|
| isid | string | The initiator portion of the session identifier specified by the initiator during login. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target_portal_group | string | The target portal group to which the session belongs. |
| target_portal_group_tag | integer | The target portal group tag of the session. |
| tsih | integer | The target session identifier handle (TSIH) of the session. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve an iSCSI session**

GET `/protocols/san/iscsi/sessions/{svm.uuid}/{tpgroup}/{tsih}`

**Introduced In:** 9.6

Retrieves an iSCSI session.

#### Related ONTAP commands

- `vserver iscsi connection show`
- `vserver iscsi session parameter show`
- `vserver iscsi session show`

#### Learn more

- DOC /protocols/san/iscsi/sessions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | The unique identifier of the SVM of the iSCSI session. |
| tpgroup | string | path | True | The target portal group of the iSCSI session. |
| tsih | integer | path | True | The target session identifying handle. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| connections | array[iscsi_connection] | The iSCSI connections that make up the iSCSI session. |
| igroups | array[igroups] | The initiator groups in which the initiator is a member. |
| initiator | initiator | The initiator that created the session. |
| isid | string | The initiator portion of the session identifier specified by the initiator during login. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| target_portal_group | string | The target portal group to which the session belongs. |
| target_portal_group_tag | integer | The target portal group tag of the session. |
| tsih | integer | The target session identifier handle (TSIH) of the session. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "connections": [
    {
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_type": "string",
      "cid": 0,
      "initiator_address": {
        "address": "10.10.10.7",
        "port": 55432
      },
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7",
          "port": 3260
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
```

```
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiator": {
    "alias": "initiator_alias1",
    "comment": "This is an iSCSI initiator for host 5",
    "name": "iqn.1992-01.example.com:string"
  },
  "isid": "61:62:63:64:65:00",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target_portal_group": "tpgroup1",
  "target_portal_group_tag": 0,
  "tsih": 0
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | An SVM with the specified UUID does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

initiator_address

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

| Name | Type | Description |
|---|---|---|
| address | string | The TCP IPv4 or IPv6 address of the initiator end of the iSCSI connection. |
| port | integer | The TCP port number of the initiator end of the iSCSI connection. |

ip

The IP information. ONTAP only supports port 3260.

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| port | integer | The TCP port number of the iSCSI access endpoint. |

interface

The network interface information for the target end of the connection.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | The IP information. ONTAP only supports port 3260.<br>• readOnly: 1<br>• Introduced in: 9.6 |
| name | string | The name of the interface. |
| uuid | string | The UUID that uniquely identifies the interface. |

iscsi_connection

An active iSCSI connection.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| authentication_type | string | The iSCSI authentication type used to establish the connection. |
| cid | integer | The identifier of the connection within the session. |
| initiator_address | initiator_address | The TCP socket information for the initiator end of the connection. This is useful for network packet debugging. |
| interface | interface | The network interface information for the target end of the connection. |

igroups

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiator

The initiator that created the session.

| Name | Type | Description |
|---|---|---|
| alias | string | The initiator alias. |
| comment | string | A comment available for use by the administrator. This is modifiable from the initiator REST endpoint directly. See PATCH /protocols/san/igroups/{igroup.uuid}/initiators/{name} . |
| name | string | The world wide unique name of the initiator. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage SAN LUN maps

### Manage SAN LUN maps

#### Overview

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between an initiator group and a LUN is many initiator groups to many LUNs.

A LUN map also configures the cluster nodes from which network paths to the LUN are advertised via the SAN protocols as part of the Selective LUN Map (SLM) functionality of ONTAP. These nodes are referred to as the reporting nodes of a LUN map. For further information, see DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes .

The LUN map REST API allows you to create, delete, and discover LUN maps, and manage the reporting nodes of a LUN map.

#### Examples

### Creating a LUN map

```
# The API:
POST /api/protocols/san/lun-maps

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/lun-maps' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "igroup": { "name":
"igroup1" }, "lun": { "name": "/vol/vol1/lun1" } }'
```

### Retrieving all of the LUN maps

```
# The API:
GET /api/protocols/san/lun-maps

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/lun-maps' -H 'Accept:
application/hal+json'

# The response:
```

```
{
"records": [
  {
    "svm": {
      "uuid": "03157e81-24c5-11e9-9ec1-005056bba643",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/03157e81-24c5-11e9-9ec1-005056bba643"
        }
      }
    },
    "lun": {
      "uuid": "a60d9862-9bee-49a6-8162-20d2421bb1a6",
      "name": "/vol/vol1/lun1",
      "_links": {
        "self": {
          "href": "/api/storage/luns/a60d9862-9bee-49a6-8162-20d2421bb1a6"
        }
      }
    },
    "igroup": {
      "uuid": "40d98b2c-24c5-11e9-9ec1-005056bba643",
      "name": "ig1",
      "_links": {
        "self": {
          "href": "/api/protocols/san/igroups/40d98b2c-24c5-11e9-9ec1-
005056bba643"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-
20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/san/lun-maps"
  }
}
}
```

**Retrieving a specific LUN map**

```
# The API:
GET /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/lun-maps/a60d9862-9bee-
49a6-8162-20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643' -H 'Accept:
application/hal+json'

# The response:
{
"svm": {
  "uuid": "03157e81-24c5-11e9-9ec1-005056bba643",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/03157e81-24c5-11e9-9ec1-005056bba643"
    }
  }
},
"lun": {
  "uuid": "a60d9862-9bee-49a6-8162-20d2421bb1a6",
  "name": "/vol/vol1/lun1",
  "node": {
    "uuid": "7d8607ea-24c1-11e9-9ec1-005056bba643",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/7d8607ea-24c1-11e9-9ec1-005056bba643"
      }
    }
  },
  "smbc": {
    "replicated": false
  },
  "_links": {
    "self": {
      "href": "/api/storage/luns/a60d9862-9bee-49a6-8162-20d2421bb1a6"
    }
  }
},
"igroup": {
  "uuid": "40d98b2c-24c5-11e9-9ec1-005056bba643",
  "name": "ig1",
  "os_type": "linux",
```

```
  "protocol": "mixed",
  "replicated": false,
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/40d98b2c-24c5-11e9-9ec1-
005056bba643"
    }
  }
},
"logical_unit_number": 0,
"reporting_nodes": [
  {
    "uuid": "11a465f5-2ac0-11eb-a303-005056bb1e81",
    "name": "node1",
    "_links": {
      "node": {
        "href": "/cluster/nodes/11a465f5-2ac0-11eb-a303-005056bb1e81"
      },
      "self": {
        "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-
20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643/reporting-
nodes/11a465f5-2ac0-11eb-a303-005056bb1e81"
      }
    }
  },
  {
    "uuid": "6c7cb50f-2abf-11eb-9840-005056bbd490",
    "name": "node2",
    "_links": {
      "node": {
        "href": "/cluster/nodes/6c7cb50f-2abf-11eb-9840-005056bbd490"
      },
      "self": {
        "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-
20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643/reporting-
nodes/6c7cb50f-2abf-11eb-9840-005056bbd490"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-
20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643"
  }
}
```

```
}
```

### Deleting a LUN map

```
# The API:
DELETE /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/lun-maps/a60d9862-
9bee-49a6-8162-20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643' -H
'Accept: application/hal+json'
```

### Retrieve LUN maps

GET `/protocols/san/lun-maps`

**Introduced In:** 9.6

Retrieves LUN maps.

#### Related ONTAP commands

- `lun mapping show`

#### Learn more

- DOC /protocols/san/lun-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.name | string | query | False | Filter by lun.name |
| lun.smbc.replicated | boolean | query | False | Filter by lun.smbc.replicated<br>• Introduced in: 9.15 |
| lun.node.uuid | string | query | False | Filter by lun.node.uuid |
| lun.node.name | string | query | False | Filter by lun.node.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| lun.uuid | string | query | False | Filter by lun.uuid |
| reporting_nodes.name | string | query | False | Filter by reporting_nodes.name<br>• Introduced in: 9.10 |
| reporting_nodes.uuid | string | query | False | Filter by reporting_nodes.uuid<br>• Introduced in: 9.10 |
| igroup.replicated | boolean | query | False | Filter by igroup.replicated<br>• Introduced in: 9.15 |
| igroup.uuid | string | query | False | Filter by igroup.uuid |
| igroup.protocol | string | query | False | Filter by igroup.protocol |
| igroup.os_type | string | query | False | Filter by igroup.os_type |
| igroup.name | string | query | False | Filter by igroup.name |
| igroup.initiators | string | query | False | Filter by igroup.initiators |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| logical_unit_number | integer | query | False | Filter by logical_unit_number<br>• Max value: 4095<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |

| Name | Type | Description |
| --- | --- | --- |
| records | array[lun_map] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "initiators": [
          "iqn.1998-01.com.corp.iscsi:name1"
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
      },
      "logical_unit_number": 1,
      "lun": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/volume1/qtree1/lun1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
```

```
            },
            "name": "node1",
            "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "reporting_nodes": [
          {
            "_links": {
              "node": {
                "href": "/api/resourcelink"
              },
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
          }
        ],
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| initiators | array[string] | The initiators that are members of the initiator group. |
| name | string | The name of the initiator group. Valid in POST. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. |
| replicated | boolean | This property reports if the initiator group is replicated. |
| uuid | string | The unique identifier of the initiator group. Valid in POST. |

node

The LUN node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the LUN's node. |
| uuid | string | The unique identifier of the LUN node. |

smbc

"Properties related to SM-BC replication."

| Name | Type | Description |
|---|---|---|
| replicated | boolean | This property reports if the LUN is replicated via SM-BC. |

lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the LUN. Valid in POST.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | The LUN node. |
| smbc | smbc | "Properties related to SM-BC replication." |
| uuid | string | The unique identifier of the LUN. Valid in POST. |

_links

| Name | Type | Description |
| --- | --- | --- |
| node | href | |
| self | href | |

reporting_nodes

A cluster node from which network paths to the LUN are advertised by ONTAP via the SAN protocols.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

lun_map

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.<br>• example: 1<br>• Max value: 4095<br>• Min value: 0<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| lun | lun | The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both. |

| Name | Type | Description |
| --- | --- | --- |
| reporting_nodes | array[reporting_nodes] | The cluster nodes from which network paths to the mapped LUNs are advertised via the SAN protocols as part of the Selective LUN Map (SLM) feature of ONTAP.<br><br>When a LUN map is created, the cluster node hosting the LUN and its high availability (HA) partner are set as the default reporting node. In POST, the property `additional_reporting_node` may be used to add an additional node and its HA partner.<br><br>For further information, see DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes .<br><br>• readOnly: 1<br>• Introduced in: 9.10 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a LUN map

POST `/protocols/san/lun-maps`

**Introduced In:** 9.6

Creates a LUN map.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the LUN map.
- `igroup.uuid` or `igroup.name` - Existing initiator group to map to the specified LUN.
- `lun.uuid` or `lun.name` - Existing LUN to map to the specified initiator group.

#### Default property values

If not specified in POST, the following default property values are assigned.

- `logical_unit_number` - If no value is provided, ONTAP assigns the lowest available value.

#### Related ONTAP commands

- `lun mapping create`

#### Learn more

- DOC /protocols/san/lun-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| additional_reporting_node.uuid | string | query | False | The unique identifier of an ONTAP cluster node to add to the default reporting nodes for the LUN map. The HA partner for the node is also added.<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| additional_reporting_node.name | string | query | False | The name of an ONTAP cluster node to add to the default reporting nodes for the LUN map. The HA partner for the node is also added.<br>• Introduced in: 9.10 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| igroup | igroup | The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.<br>• example: 1<br>• Max value: 4095<br>• Min value: 0<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| lun | lun | The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both. |
| reporting_nodes | array[reporting_nodes] | The cluster nodes from which network paths to the mapped LUNs are advertised via the SAN protocols as part of the Selective LUN Map (SLM) feature of ONTAP.<br><br>When a LUN map is created, the cluster node hosting the LUN and its high availability (HA) partner are set as the default reporting node. In POST, the property `additional_reporting_node` may be used to add an additional node and its HA partner.<br><br>For further information, see DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes .<br><br>• readOnly: 1<br>• Introduced in: 9.10 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "igroup": {
    "initiators": [
      "iqn.1998-01.com.corp.iscsi:name1"
    ],
    "name": "igroup1",
    "os_type": "string",
    "protocol": "string",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": 1,
  "lun": {
    "name": "/vol/volume1/qtree1/lun1",
    "node": {
      "name": "node1",
      "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "reporting_nodes": [
    {
      "name": "node1",
      "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[lun_map] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "igroup": {
        "initiators": [
          "iqn.1998-01.com.corp.iscsi:name1"
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
      },
      "logical_unit_number": 1,
      "lun": {
        "name": "/vol/volume1/qtree1/lun1",
        "node": {
          "name": "node1",
          "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "reporting_nodes": [
        {
          "name": "node1",
          "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254201 | The specified LUN does not exist. |
| 1254207 | The LUN is already mapped to the same initiator group. |
| 2621462 | The specified SVM does not exist or is not accessible to the caller. |
| 2621706 | Both the SVM UUID and SVM name were supplied, but don't refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374053 | The LUN is the destination of an ongoing restore operation and is inaccessible for I/O and management. Wait for the restore to complete and try the command again. |
| 5374127 | The specified LUN name is invalid. |
| 5374238 | The operation is not allowed on a LUN in a snapshot. |
| 5374316 | A LUN move operation is in progress on the source LUN. |
| 5374329 | A LUN of class `vvol` cannot be mapped. |
| 5374573 | A node has no interface configured with the iSCSI or Fibre Channel protocols for the specified SVM. |
| 5374574 | Multiple nodes have no interface configured with the iSCSI or Fibre Channel protocols for the specified SVM. |
| 5374581 | A node has no interface configured with the iSCSI protocol for the specified SVM. |
| 5374582 | Multiple nodes have no interface configured with the iSCSI protocol for the specified SVM. |
| 5374583 | A node has no interface configured with the Fibre Channel protocol for the specified SVM. |
| 5374584 | Multiple nodes have no interface configured with the Fibre Channel protocol for the specified SVM. |
| 5374652 | An error was reported by the peer cluster while creating a replicated LUN map. The specific error will be included as a nested error. |

| Error Code | Description |
| --- | --- |
| 5374901 | Either `lun.uuid` or `lun.name` must be provided to create a LUN map. |
| 5374902 | Either `igroup.uuid` or `igroup.name` must be provided to create a LUN map. |
| 5374920 | The cluster node specified in additional_reporting_node does not exist. |
| 5374921 | The cluster node name and UUID specified in additional_reporting_node do not refer to the same cluster node. |
| 5376356 | The LUN is part of a SnapMirror active sync relationship that has a different SVM peering relationship than the igroup. |
| 5376461 | The specified LUN name is invalid. |
| 5376462 | The specified LUN name is too long. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

| Name | Type | Description |
|---|---|---|
| initiators | array[string] | The initiators that are members of the initiator group. |
| name | string | The name of the initiator group. Valid in POST. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. |
| replicated | boolean | This property reports if the initiator group is replicated. |
| uuid | string | The unique identifier of the initiator group. Valid in POST. |

node

The LUN node.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the LUN's node. |
| uuid | string | The unique identifier of the LUN node. |

smbc

"Properties related to SM-BC replication."

| Name | Type | Description |
|---|---|---|
| replicated | boolean | This property reports if the LUN is replicated via SM-BC. |

lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the LUN. Valid in POST.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | The LUN node. |
| smbc | smbc | "Properties related to SM-BC replication." |
| uuid | string | The unique identifier of the LUN. Valid in POST. |

reporting_nodes

A cluster node from which network paths to the LUN are advertised by ONTAP via the SAN protocols.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

lun_map

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both. |

| Name | Type | Description |
| --- | --- | --- |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.<br>• example: 1<br>• Max value: 4095<br>• Min value: 0<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| lun | lun | The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both. |
| reporting_nodes | array[reporting_nodes] | The cluster nodes from which network paths to the mapped LUNs are advertised via the SAN protocols as part of the Selective LUN Map (SLM) feature of ONTAP.<br>When a LUN map is created, the cluster node hosting the LUN and its high availability (HA) partner are set as the default reporting node. In POST, the property `additional_reporting_node` may be used to add an additional node and its HA partner.<br>For further information, see DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes .<br>• readOnly: 1<br>• Introduced in: 9.10 |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a LUN map

DELETE `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}`

**Introduced In:** 9.6

Deletes a LUN map.

#### Related ONTAP commands

- `lun mapping delete`

#### Learn more

- DOC /protocols/san/lun-maps

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| lun.uuid | string | path | True | |
| igroup.uuid | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN does not exist or is not accessible to the caller. |
| 5374878 | The specified initiator group does not exist, is not accessible to the caller, or is not in the same SVM as the specified LUN. |
| 5374922 | The specified LUN map does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a LUN map

GET `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}`

**Introduced In:** 9.6

Retrieves a LUN map.

#### Related ONTAP commands

- `lun mapping show`

#### Learn more

- DOC /protocols/san/lun-maps

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| igroup.uuid | string | path | True | The unique identifier of the igroup. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.<br>• example: 1<br>• Max value: 4095<br>• Min value: 0<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |
| lun | lun | The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both. |

| Name | Type | Description |
| --- | --- | --- |
| reporting_nodes | array[reporting_nodes] | The cluster nodes from which network paths to the mapped LUNs are advertised via the SAN protocols as part of the Selective LUN Map (SLM) feature of ONTAP.<br><br>When a LUN map is created, the cluster node hosting the LUN and its high availability (HA) partner are set as the default reporting node. In POST, the property `additional_reporting_node` may be used to add an additional node and its HA partner.<br><br>For further information, see DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes .<br><br>• readOnly: 1<br>• Introduced in: 9.10 |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "initiators": [
      "iqn.1998-01.com.corp.iscsi:name1"
    ],
    "name": "igroup1",
    "os_type": "string",
    "protocol": "string",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": 1,
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/qtree1/lun1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "reporting_nodes": [
    {
      "_links": {
        "node": {
          "href": "/api/resourcelink"
```

```
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374852 | The initiator group does not exist or is not accessible to the caller. |
| 5374875 | The LUN does not exist or is not accessible to the caller. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| initiators | array[string] | The initiators that are members of the initiator group. |
| name | string | The name of the initiator group. Valid in POST. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. |
| replicated | boolean | This property reports if the initiator group is replicated. |
| uuid | string | The unique identifier of the initiator group. Valid in POST. |

node

The LUN node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN's node. |
| uuid | string | The unique identifier of the LUN node. |

smbc

"Properties related to SM-BC replication."

| Name | Type | Description |
| --- | --- | --- |
| replicated | boolean | This property reports if the LUN is replicated via SM-BC. |

lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the LUN. Valid in POST.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| node | node | The LUN node. |
| smbc | smbc | "Properties related to SM-BC replication." |
| uuid | string | The unique identifier of the LUN. Valid in POST. |

_links

| Name | Type | Description |
| --- | --- | --- |
| node | href | |
| self | href | |

reporting_nodes

A cluster node from which network paths to the LUN are advertised by ONTAP via the SAN protocols.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Add, remove, or discover LUN map reporting nodes

### Add, remove, or discover LUN map reporting nodes

#### Overview

A LUN map's reporting nodes are the cluster nodes from which network paths to a mapped LUN are advertised via the SAN protocols as part of the Selective LUN Map (SLM) feature of ONTAP. SLM reduces the number of paths from the host to a mapped LUN and enables management of a single initiator group (igroup) per host.

If there are no reporting nodes in a LUN map, network paths to all cluster nodes having the appropriate network interfaces (LIFs) in the SVM are advertised. This is not a typical configuration and is reserved for limited specific use cases. Note that having no reporting nodes in a LUN map differs subtly from having all reporting nodes in the LUN map. If a LUN map has an empty reporting nodes list and a new node is added to the cluster, a path to the new node will also be advertised. If a LUN map has all cluster nodes in its reporting nodes list and a new node is added to the cluster, a path to the new node is not advertised unless the LUN map's reporting nodes are explicitly updated to include the new node.

If portsets are used to further restrict access for initiators to specific LIFs, the mapped LUN will be accessible only via the LIFs in the portset that are on the reporting nodes of the LUN map.

When a LUN map is created, the cluster node hosting the mapped LUN and its high availability (HA) partner are set as the initial reporting nodes.

Before moving a mapped LUN or a volume containing mapped LUNs to another HA pair within the same cluster, the destination node should be added to the LUN map's reporting nodes. This ensures that active, optimized LUN paths are maintained. After moving a mapped LUN or a volume containing mapped LUNs to another HA pair within the same cluster, the cluster node that previously hosted the mapped LUN should be removed from the LUN map's reporting node. Further details for this workflow may be found in the ONTAP SAN Administration documentation - see `Modifying the SLM reporting-nodes list`.

#### Examples

### Adding a node to a LUN map

This example adds a cluster node, and its high availability (HA) partner cluster node, to a LUN map's reporting nodes.

```
# The API:
POST /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/lun-maps/c8ad35ce-ef03-
4577-88f0-9ad150ac6bbf/b10a8165-8346-11eb-ab8e-005056bbb402/reporting-
nodes' -H 'Accept: application/hal+json' -d '{ "name": "node2" }'
```

---

#### Removing a node from a LUN map

This example removes a cluster node, and its high availability (HA) partner cluster node, from a LUN map's reporting nodes.

```
# The API:
DELETE /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-
nodes/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/lun-maps/c8ad35ce-
ef03-4577-88f0-9ad150ac6bbf/b10a8165-8346-11eb-ab8e-
005056bbb402/reporting-nodes/6d2cd7d5-493a-daf8-9ae1-219e4ad6f77d' -H
'Accept: application/hal+json'
```

#### Retrieve LUN map reporting nodes

GET `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes`

**Introduced In:** 9.10

Retrieves LUN map reporting nodes.

##### Related ONTAP commands

- `lun mapping show`

##### Learn more

- DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes

##### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the igroup. |
| name | string | query | False | Filter by name |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[lun_map_reporting_node] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "node": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "lun": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "node1",
      "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN does not exist or is not accessible to the caller. |
| 5374878 | The specified initiator group does not exist, is not accessible to the caller, or is not in the same SVM as the specified LUN. |
| 5374922 | The specified LUN map does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| node | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

The initiator group of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the initiator group. |

lun

The LUN of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the LUN. |

lun_map_reporting_node

A cluster node from which network paths to the LUN are advertised by ONTAP via the SAN protocols.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | The initiator group of the LUN map of the reporting node. |
| lun | lun | The LUN of the LUN map of the reporting node. |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Add a reporting node and its HA partner to a LUN map**

POST `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes`

**Introduced In:** 9.10

Adds a reporting node and its HA partner to a LUN map.

#### Required properties

- `uuid` or `name` - A cluster node to add.

#### Related ONTAP commands

- `lun mapping add-reporting-nodes`

#### Learn more

- DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| igroup.uuid | string | path | True | The unique identifier of the igroup. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | The initiator group of the LUN map of the reporting node. |
| lun | lun | The LUN of the LUN map of the reporting node. |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

### Example request

```
{
  "igroup": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "lun": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "node1",
  "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[lun_map_reporting_node] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "igroup": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "lun": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "node1",
      "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN does not exist or is not accessible to the caller. |
| 5374878 | The specified initiator group does not exist, is not accessible to the caller, or is not in the same SVM as the specified LUN. |
| 5374920 | The specified cluster node does not exist. |
| 5374921 | The specified cluster node name and UUID do not refer to the same cluster node. |
| 5374922 | The specified LUN map does not exist. |
| 5374923 | A cluster node `uuid` or `name` must be specified to add a reporting node. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

igroup

The initiator group of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the initiator group. |

lun

The LUN of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the LUN. |

lun_map_reporting_node

A cluster node from which network paths to the LUN are advertised by ONTAP via the SAN protocols.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | The initiator group of the LUN map of the reporting node. |
| lun | lun | The LUN of the LUN map of the reporting node. |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Remove a reporting node and its HA partner from a LUN map

DELETE `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes/{uuid}`

**Introduced In:** 9.10

Removes a reporting node and its HA partner from a LUN map.

#### Related ONTAP commands

- `lun mapping remove-reporting-nodes`

#### Learn more

- DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | |
| igroup.uuid | string | path | True | |
| uuid | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374608 | An attempt was made to delete the last reporting node for a LUN. The REST API requires that a LUN map must maintain at least one reporting node. |
| 5374875 | The specified LUN does not exist or is not accessible to the caller. |
| 5374878 | The specified initiator group does not exist, is not accessible to the caller, or is not in the same SVM as the specified LUN. |
| 5374922 | The specified LUN map does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a LUN map reporting node

GET `/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes/{uuid}`

**Introduced In:** 9.10

Retrieves a LUN map reporting node.

#### Related ONTAP commands

- `lun mapping show`

#### Learn more

- DOC /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}/reporting-nodes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroup.uuid | string | path | True | The unique identifier of the igroup. |
| uuid | string | path | True | The unique identifier of the node. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroup | igroup | The initiator group of the LUN map of the reporting node. |
| lun | lun | The LUN of the LUN map of the reporting node. |
| name | string | The name of the node.<br><br>Either `uuid` or `name` are required in POST. |
| uuid | string | The unique identifier of the node.<br><br>Either `uuid` or `name` are required in POST. |

#### Example response

```
{
  "_links": {
    "node": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "node1",
  "uuid": "5ac8eb9c-4e32-dbaa-57ca-fb905976f54e"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN does not exist or is not accessible to the caller. |
| 5374878 | The specified initiator group does not exist, is not accessible to the caller, or is not in the same SVM as the specified LUN. |

| Error Code | Description |
| --- | --- |
| 5374922 | The specified LUN map does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| node | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroup

The initiator group of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the initiator group. |

lun

The LUN of the LUN map of the reporting node.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the LUN. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SAN portsets

### Manage SAN portsets

#### Overview

A portset is a collection of Fibre Channel Protocol and/or iSCSI network interfaces from the portset's SVM.

Portsets are used to limit the network interfaces through which an initiator can connect to mapped LUNs. When a portset is bound to an initiator group (igroup), and the initiator group is mapped to a LUN, the initiators of the initiator group can only reach the LUN through the network interfaces in the portset.

Portsets are not recommended for new configurations. With modern releases of ONTAP, it is recommended to use multiple SVMs and initiator groups with no bound portset to load balance applications over multiple ports on a node. Selective LUN mapping will automatically limit the number of visible paths to a LUN from the client host to those required for efficient access and high availability. The REST portset API is primarily intended for legacy use.

The portset REST API allows you to create, delete, and discover portsets, and to add and remove network interfaces from portsets.

A portset can be bound to one or more initiator groups. An initiator group (igroup) can be bound to at most one portset.

When a portset is created, the `protocol` property is used to restrict member network interfaces to Fibre Channel Protocol (*fcp*), iSCSI (*iscsi*), or both (*mixed*).

Zero or more network interfaces can be supplied when the portset is created. After creation, network interfaces can be added to or removed from the portset using the `/protocols/san/portsets/{portset.uuid}/interfaces` endpoint. See DELETE /protocols/san/portsets/{portset.uuid}/interfaces/{name} for more details.

#### Examples

### Creating a portset with no network interfaces

The example portset uses the default `mixed` protocol. Note that the `return_records` query parameter is used to obtain the newly created portset in the response.

```
# The API:
POST /api/protocols/san/portsets

# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/portsets?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm" : { "name": "svm1" }, "name": "portset1"
}'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "5d7b1dfa-1ed7-11eb-8b0f-005056bb3521",
    "svm": {
      "uuid": "31484775-1e23-11eb-b2a8-005056bb3521",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/31484775-1e23-11eb-b2a8-005056bb3521"
        }
      }
    },
    "name": "portset1",
    "protocol": "mixed",
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/5d7b1dfa-1ed7-11eb-8b0f-
005056bb3521"
      }
    }
  }
]
}
```

#### Creating an iSCSI portset with two network interfaces

Note that the `return_records` query parameter is used to obtain the newly created portset in the response.

```
# The API:
POST /api/protocols/san/portsets
```

```
# The call:
curl -X POST 'https://<mgmt-
ip>/api/protocols/san/portsets?return_records=true' -H 'Accept:
application/hal+json' -d '{ "svm" : { "name": "svm1" }, "name":
"portset2", "protocol": "iscsi", "interfaces": [ { "ip": { "name": "lif1"
} }, { "ip": { "name": "lif2" } } ] }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "816c0d49-1ed7-11eb-8b0f-005056bb3521",
    "svm": {
      "uuid": "31484775-1e23-11eb-b2a8-005056bb3521",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/31484775-1e23-11eb-b2a8-005056bb3521"
        }
      }
    },
    "name": "portset2",
    "protocol": "iscsi",
    "interfaces": [
      {
        "uuid": "f37bfb01-1e2a-11eb-b2a8-005056bb3521",
        "ip": {
          "uuid": "f37bfb01-1e2a-11eb-b2a8-005056bb3521",
          "name": "lif1",
          "ip": {
            "address": "192.168.1.100"
          },
          "_links": {
            "self": {
              "href": "/api/network/ip/interfaces/f37bfb01-1e2a-11eb-b2a8-
005056bb3521"
            }
          }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/portsets/816c0d49-1ed7-11eb-8b0f-
005056bb3521/interfaces/f37bfb01-1e2a-11eb-b2a8-005056bb3521"
          }
        }
```

```
      },
      {
        "uuid": "f92178e7-1e2a-11eb-b2a8-005056bb3521",
        "ip": {
          "uuid": "f92178e7-1e2a-11eb-b2a8-005056bb3521",
          "name": "lif2",
          "ip": {
            "address": "192.168.1.101"
          },
          "_links": {
            "self": {
              "href": "/api/network/ip/interfaces/f92178e7-1e2a-11eb-b2a8-
005056bb3521"
            }
          }
        },
        "_links": {
          "self": {
            "href": "/api/protocols/san/portsets/816c0d49-1ed7-11eb-8b0f-
005056bb3521/interfaces/f92178e7-1e2a-11eb-b2a8-005056bb3521"
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/816c0d49-1ed7-11eb-8b0f-
005056bb3521"
      }
    }
  }
]
}
```


---

#### Retrieving a summary of all portsets

```
# The API:
GET /api/protocols/san/portsets

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/portsets' -H 'Accept:
application/hal+json'

# The response:
```

```
{
"records": [
  {
    "uuid": "5d7b1dfa-1ed7-11eb-8b0f-005056bb3521",
    "name": "portset1",
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/5d7b1dfa-1ed7-11eb-8b0f-
005056bb3521"
      }
    }
  },
  {
    "uuid": "816c0d49-1ed7-11eb-8b0f-005056bb3521",
    "name": "portset2",
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/816c0d49-1ed7-11eb-8b0f-
005056bb3521"
      }
    }
  },
  {
    "uuid": "b716b4d2-1ed7-11eb-8b0f-005056bb3521",
    "name": "portset3",
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/b716b4d2-1ed7-11eb-8b0f-
005056bb3521"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/protocols/san/portsets"
  }
}
}
```


#### Retrieving details for a specific portset

```
# The API:
```

```
GET /api/protocols/san/portsets/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/portsets/b716b4d2-1ed7-
11eb-8b0f-005056bb3521' -H 'Accept: application/hal+json'

# The response:
{
"uuid": "b716b4d2-1ed7-11eb-8b0f-005056bb3521",
"svm": {
  "uuid": "31484775-1e23-11eb-b2a8-005056bb3521",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/31484775-1e23-11eb-b2a8-005056bb3521"
    }
  }
},
"name": "portset3",
"protocol": "fcp",
"interfaces": [
  {
    "uuid": "164eb052-1e2b-11eb-b2a8-005056bb3521",
    "fc": {
      "uuid": "164eb052-1e2b-11eb-b2a8-005056bb3521",
      "name": "lif5",
      "wwpn": "20:01:00:50:56:bb:35:21",
      "_links": {
        "self": {
          "href": "/api/network/fc/interfaces/164eb052-1e2b-11eb-b2a8-
005056bb3521"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/b716b4d2-1ed7-11eb-8b0f-
005056bb3521/interfaces/164eb052-1e2b-11eb-b2a8-005056bb3521"
      }
    }
  },
  {
    "uuid": "197ba2b7-1e2b-11eb-b2a8-005056bb3521",
    "fc": {
      "uuid": "197ba2b7-1e2b-11eb-b2a8-005056bb3521",
      "name": "lif6",
```

```
      "wwpn": "20:02:00:50:56:bb:35:21",
      "_links": {
        "self": {
          "href": "/api/network/fc/interfaces/197ba2b7-1e2b-11eb-b2a8-
005056bb3521"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/portsets/b716b4d2-1ed7-11eb-8b0f-
005056bb3521/interfaces/197ba2b7-1e2b-11eb-b2a8-005056bb3521"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/portsets/b716b4d2-1ed7-11eb-8b0f-
005056bb3521"
  }
}
}
```


### Deleting a portset

```
# The API:
DELETE /api/protocols/san/portsets/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/portsets/b716b4d2-
1ed7-11eb-8b0f-005056bb3521' -H 'Accept: application/hal+json'
```

### Adding a network interface to a portset

```
# The API:
POST /api/protocols/san/portsets/{uuid}/interfaces

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/portsets/5d7b1dfa-1ed7-
11eb-8b0f-005056bb3521/interfaces' -H 'Accept: application/hal+json' -d '{
"fc": { "name": "lif4" } }'
}
```

---

#### Adding multiple network interfaces to a portset in a single call

```
# The API:
POST /api/protocols/san/portsets/{uuid}/interfaces

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/portsets/5d7b1dfa-1ed7-
11eb-8b0f-005056bb3521/interfaces' -H 'Accept: application/hal+json' -d '{
"records": [ { "ip": { "name": "lif1" } },  { "ip": { "name": "lif2" } },
{ "fc": { "name": "lif5" } } ] }'
```

---

#### Removing a network interface from a portset

```
# The API:
DELETE /api/protocols/san/portsets/{uuid}/interfaces/f92178e7-1e2a-11eb-
b2a8-005056bb3521

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/portsets/5d7b1dfa-
1ed7-11eb-8b0f-005056bb3521/interfaces/f92178e7-1e2a-11eb-b2a8-
005056bb3521' -H 'Accept: application/hal+json'
```

#### Retrieve portsets

GET `/protocols/san/portsets`

**Introduced In:** 9.9

Retrieves portsets.

#### Related ONTAP commands

- `lun portset show`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| name | string | query | False | Filter by name<br>• maxLength: 96<br>• minLength: 1 |
| protocol | string | query | False | Filter by protocol |
| uuid | string | query | False | Filter by uuid |
| interfaces.fc.name | string | query | False | Filter by interfaces.fc.name |
| interfaces.fc.uuid | string | query | False | Filter by interfaces.fc.uuid |
| interfaces.fc.wwpn | string | query | False | Filter by interfaces.fc.wwpn |
| interfaces.ip.ip.address | string | query | False | Filter by interfaces.ip.ip.address |
| interfaces.ip.name | string | query | False | Filter by interfaces.ip.name |
| interfaces.ip.uuid | string | query | False | Filter by interfaces.ip.uuid |
| interfaces.uuid | string | query | False | Filter by interfaces.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| igroups.name | string | query | False | Filter by igroups.name<br>• maxLength: 96<br>• minLength: 1 |
| igroups.uuid | string | query | False | Filter by igroups.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[portset] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroups": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "fc": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "fc_lif1",
            "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
            "wwpn": "20:00:00:50:56:b4:13:a8"
```

```
            },
            "ip": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "ip": {
                "address": "10.10.10.7"
              },
              "name": "lif1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "name": "portset1",
        "protocol": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroups

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interfaces

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

portset

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| igroups | array[igroups] | An array initiator groups to which the portset is bound. |
| interfaces | array[interfaces] | An array of network interfaces that are members of the portset. These are the only network interfaces through which the initiators of a bound igroup can access mapped LUNs.<br><br>Zero or more network interfaces can be supplied when the portset is created. After creation, network interfaces can be added or removed from the portset using the `/protocols/san/portsets/{portset.uuid}/interfaces` endpoint. See DELETE /protocols/san/portsets/{portset.uuid}/interfaces/{uuid} for more details. |
| name | string | The name of the portset. Required in POST.<br><br>The name of a portset cannot be changed after creation. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the portset. This restricts the type of network interfaces that can be added to the portset. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of a portset cannot be changed after creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier for a portset. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a portset**

POST `/protocols/san/portsets`

**Introduced In:** 9.9

Creates a portset.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the portset.
- `name` - Name of the portset.

#### Recommended optional properties

- `protocol` - The network protocol of the interfaces in the portset.
- `interfaces` - Network interfaces to include in the portset. This property can be used to create the portset and populate it with network interfaces in a single request.

#### Default property values

If not specified in POST, the following default property values are assigned.

- `protocol` - *mixed* - Data protocol of the portset's network interfaces.

#### Related ONTAP commands

- `lun portset create`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| igroups | array[igroups] | An array initiator groups to which the portset is bound. |

| Name | Type | Description |
| --- | --- | --- |
| interfaces | array[interfaces] | An array of network interfaces that are members of the portset. These are the only network interfaces through which the initiators of a bound igroup can access mapped LUNs.<br><br>Zero or more network interfaces can be supplied when the portset is created. After creation, network interfaces can be added or removed from the portset using the `/protocols/san/portsets/{portset.uuid}/interfaces` endpoint. See DELETE /protocols/san/portsets/{portset.uuid}/interfaces/{uuid} for more details. |
| name | string | The name of the portset. Required in POST.<br><br>The name of a portset cannot be changed after creation. |
| protocol | string | The protocols supported by the portset. This restricts the type of network interfaces that can be added to the portset. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of a portset cannot be changed after creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier for a portset. |

### Example request

```
{
  "igroups": [
    {
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "interfaces": [
    {
      "fc": {
        "name": "fc_lif1",
        "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "ip": {
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "portset1",
  "protocol": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[portset] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "igroups": [
        {
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "interfaces": [
        {
          "fc": {
            "name": "fc_lif1",
            "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
            "wwpn": "20:00:00:50:56:b4:13:a8"
          },
          "ip": {
            "name": "lif1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "name": "portset1",
      "protocol": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254259 | A portset with the same name already exists in the SVM. |
| 2621462 | The specified SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5373958 | The specified portset name contains invalid characters. |
| 5373982 | An invalid WWN was specified. The length is incorrect. |
| 5373983 | An invalid WWN was specified. The format is incorrect. |
| 5374905 | An invalid interfaces array element was specified. |
| 5374906 | A specified network interface was not found. |
| 5374907 | The specified network interface UUID and name don't identify the same network interface. |
| 5374914 | An attempt was made to add a network interface of an incompatible protocol to a portset. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

igroups

| Name | Type | Description |
|---|---|---|
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID that uniquely identifies the interface. |

interfaces

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

portset

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | An array initiator groups to which the portset is bound. |

| Name | Type | Description |
| --- | --- | --- |
| interfaces | array[interfaces] | An array of network interfaces that are members of the portset. These are the only network interfaces through which the initiators of a bound igroup can access mapped LUNs.<br><br>Zero or more network interfaces can be supplied when the portset is created. After creation, network interfaces can be added or removed from the portset using the `/protocols/san/portsets/{portset.uuid}/interfaces` endpoint. See DELETE /protocols/san/portsets/{portset.uuid}/interfaces/{uuid} for more details. |
| name | string | The name of the portset. Required in POST.<br><br>The name of a portset cannot be changed after creation. |
| protocol | string | The protocols supported by the portset. This restricts the type of network interfaces that can be added to the portset. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of a portset cannot be changed after creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier for a portset. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve portset network interfaces

GET `/protocols/san/portsets/{portset.uuid}/interfaces`

**Introduced In:** 9.9

Retrieves interfaces of a portset.

#### Related ONTAP commands

- `lun portset show`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| portset.uuid | string | path | True | The unique identifier of the portset. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[portset_interface] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "fc": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "fc_lif1",
        "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "ip": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "portset": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
          },
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374908 | The portset specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

portset

The portset in which the network interface is found.

Note that this does not mean that the network interface cannot also be found in other portsets.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The unique identifier of the portset. |

records

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

portset_interface

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| fc | fc | An FC interface. |

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| portset | portset | The portset in which the network interface is found.<br><br>Note that this does not mean that the network interface cannot also be found in other portsets. |
| uuid | string | The unique identifier of the network interface. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add portset network interfaces

POST `/protocols/san/portsets/{portset.uuid}/interfaces`

**Introduced In:** 9.9

Adds one or more interfaces to a portset.

#### Required properties

- `fc`, `ip` or `records` - Network interface(s) to add to the portset.

#### Related ONTAP commands

- `lun portset add`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| portset.uuid | string | path | True | The unique identifier of the portset. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| portset | portset | The portset in which the network interface is found.<br><br>Note that this does not mean that the network interface cannot also be found in other portsets. |
| records | array[records] | An array of network interfaces specified to add multiple interfaces to a portset in a single API call. Valid in POST only and not allowed when the `fc` or `ip` property is used. |
| uuid | string | The unique identifier of the network interface. |

### Example request

```
{
  "fc": {
    "name": "fc_lif1",
    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "ip": {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "portset": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "records": [
    {
      "fc": {
        "name": "fc_lif1",
        "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "ip": {
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[portset_interface] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "fc": {
        "name": "fc_lif1",
        "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "ip": {
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "portset": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "records": [
        {
          "fc": {
            "name": "fc_lif1",
            "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
            "wwpn": "20:00:00:50:56:b4:13:a8"
          },
          "ip": {
            "name": "lif1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374906 | A specified network interface was not found. |
| 5374907 | The specified network interface UUID and name don't identify the same network interface. |
| 5374909 | An invalid combination of network interface properties was specified. |
| 5374910 | An incomplete set of network interface properties was specified. |
| 5374914 | An attempt was made to add a network interface of an incompatible protocol to a portset. |
| 5374915 | An attempt was made to add a duplicate network interface to a portset. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

portset

The portset in which the network interface is found.

Note that this does not mean that the network interface cannot also be found in other portsets.

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the portset. |

records

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

portset_interface

| Name | Type | Description |
| --- | --- | --- |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| portset | portset | The portset in which the network interface is found.<br><br>Note that this does not mean that the network interface cannot also be found in other portsets. |
| records | array[records] | An array of network interfaces specified to add multiple interfaces to a portset in a single API call. Valid in POST only and not allowed when the `fc` or `ip` property is used. |
| uuid | string | The unique identifier of the network interface. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete portset network interfaces

DELETE `/protocols/san/portsets/{portset.uuid}/interfaces/{uuid}`

**Introduced In:** 9.9

Deletes a network interface from a portset.

#### Related ONTAP commands

- `lun portset remove`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| portset.uuid | string | path | True | The unique identifier of the portset. |
| uuid | string | path | True | The unique identifier of the network interface in the portset. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374906 | A specified network interface was not found. |
| 5374908 | The portset specified in the URI does not exist. |
| 5374916 | The specified network interface is not in the portset. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve specific portset network interfaces

GET `/protocols/san/portsets/{portset.uuid}/interfaces/{uuid}`

**Introduced In:** 9.9

Retrieves a network interface of a portset.

#### Related ONTAP commands

- `lun portset show`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| portset.uuid | string | path | True | The unique identifier of the portset. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the network interface in the portset. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| portset | portset | The portset in which the network interface is found.<br><br>Note that this does not mean that the network interface cannot also be found in other portsets. |
| uuid | string | The unique identifier of the network interface. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "fc": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "fc_lif1",
    "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "ip": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "portset": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The interface is not a member of the portset. |
| 5374908 | The portset specified in the URI does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

portset

The portset in which the network interface is found.

Note that this does not mean that the network interface cannot also be found in other portsets.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| uuid | string | The unique identifier of the portset. |

records

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a portset

DELETE `/protocols/san/portsets/{uuid}`

**Introduced In:** 9.9

Deletes a portset.

#### Related ONTAP commands

- `lun portset delete`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the portset. |
| allow_delete_while_bound | boolean | query | False | Allows deletion of a portset that is bound to an igroup.<br><br>Deleting a portset can expand the set of LIFs through which a LUN is available.<br><br>• Default value: |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374063 | An attempt was made to delete a portset that is bound to an initiator group without using query parameter `allow_delete_while_bound`. |
| 5374908 | The portset does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a portset

GET `/protocols/san/portsets/{uuid}`

**Introduced In:** 9.9

Retrieves a portset.

#### Related ONTAP commands

- `lun portset show`

#### Learn more

- DOC /protocols/san/portsets

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the portset. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| igroups | array[igroups] | An array initiator groups to which the portset is bound. |
| interfaces | array[interfaces] | An array of network interfaces that are members of the portset. These are the only network interfaces through which the initiators of a bound igroup can access mapped LUNs.<br><br>Zero or more network interfaces can be supplied when the portset is created. After creation, network interfaces can be added or removed from the portset using the `/protocols/san/portsets/{portset.uuid}/interfaces` endpoint. See DELETE /protocols/san/portsets/{portset.uuid}/interfaces/{uuid} for more details. |
| name | string | The name of the portset. Required in POST.<br><br>The name of a portset cannot be changed after creation. |
| protocol | string | The protocols supported by the portset. This restricts the type of network interfaces that can be added to the portset. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of a portset cannot be changed after creation. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The unique identifier for a portset. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "fc": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "fc_lif1",
        "uuid": "3a09ab42-4da1-32cf-9d35-3385a6101a0b",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "ip": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
```

```
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "name": "portset1",
    "protocol": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The portset does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

igroups

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

fc

An FC interface.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC interface. |
| uuid | string | The unique identifier of the FC interface. |
| wwpn | string | The WWPN of the FC interface. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

ip

A network interface. Either UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

interfaces

A container for either a Fibre Channel network interface or an IP network interface. On POST `fc` and `ip` are mutually exclusive.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| fc | fc | An FC interface. |
| ip | ip | A network interface. Either UUID or name may be supplied on input. |
| uuid | string | The unique identifier of the network interface. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SAN vVol bindings

### Manage SAN vVol bindings

#### Overview

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

Class `protocol_endpoint` and `vvol` LUNs support many-to-many vVol bindings. A LUN of one class can be bound to zero or more LUNs of the opposite class.

The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed.

The vVol binding REST API allows you to create, delete, and discover vVol bindings.

#### Examples

### Creating a vVol binding

```
# The API:
POST /api/protocols/san/vvol-bindings

# The call:
curl -X POST 'https://<mgmt-ip>/api/protocols/san/vvol-bindings' -H
'Accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"protocol_endpoint": { "name": "/vol/vol1/pe1" }, "vvol" : { "name":
"/vol/vol1/vvol1" } }'
```

#### Retrieving all vVol bindings

```
# The API:
GET /api/protocols/san/vvol-bindings

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/vvol-bindings' -H
'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "protocol_endpoint": {
      "uuid": "2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4",
      "name": "/vol/vol1/pe1",
      "_links": {
        "self": {
          "href": "/api/storage/luns/2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4"
        }
      }
    },
    "vvol": {
      "uuid": "28c02623-42fa-4f5f-a984-a02044bfc005",
      "name": "/vol/vol1/vvol1",
      "_links": {
        "self": {
          "href": "/api/storage/luns/28c02623-42fa-4f5f-a984-a02044bfc005"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/vvol-bindings/2aab57f3-dc5d-491e-80d2-
```

```
15c7ed5dd5c4/28c02623-42fa-4f5f-a984-a02044bfc005"
      }
    }
  },
  {
    "protocol_endpoint": {
      "uuid": "2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4",
      "name": "/vol/vol1/pe1",
      "_links": {
        "self": {
          "href": "/api/storage/luns/2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4"
        }
      }
    },
    "vvol": {
      "uuid": "a8d4ba93-918f-40ad-a1e4-4d7b244bdcdf",
      "name": "/vol/vol1/vvol2",
      "_links": {
        "self": {
          "href": "/api/storage/luns/a8d4ba93-918f-40ad-a1e4-4d7b244bdcdf"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/vvol-bindings/2aab57f3-dc5d-491e-80d2-
15c7ed5dd5c4/a8d4ba93-918f-40ad-a1e4-4d7b244bdcdf"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/vvol-bindings"
  }
}
}
```


#### Retrieving a specific vVol binding

```
# The API:
GET /api/protocols/san/vvol-bindings/{protocol_endpoint.uuid}/{vvol.uuid}
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/vvol-bindings/2aab57f3-
dc5d-491e-80d2-15c7ed5dd5c4/28c02623-42fa-4f5f-a984-a02044bfc005' -H
'Accept: application/hal+json'

# The response:
{
"protocol_endpoint": {
  "uuid": "2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4",
  "name": "/vol/vol1/pe1",
  "_links": {
    "self": {
      "href": "/api/storage/luns/2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4"
    }
  }
},
"vvol": {
  "uuid": "28c02623-42fa-4f5f-a984-a02044bfc005",
  "name": "/vol/vol1/vvol1",
  "_links": {
    "self": {
      "href": "/api/storage/luns/28c02623-42fa-4f5f-a984-a02044bfc005"
    }
  }
},
"svm": {
  "uuid": "bf295ccc-a6bb-11eb-93e8-005056bb470f",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/bf295ccc-a6bb-11eb-93e8-005056bb470f"
    }
  }
},
"id": 2411392,
"is_optimal": true,
"count": 1,
"_links": {
  "self": {
    "href": "/api/protocols/san/vvol-bindings/2aab57f3-dc5d-491e-80d2-
15c7ed5dd5c4/28c02623-42fa-4f5f-a984-a02044bfc005"
  }
}
}
```

### Deleting a vVol binding

```
# The API:
DELETE /api/protocols/san/vvol-
bindings/{protocol_endpoint.uuid}/{vvol.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/vvol-
bindings/2aab57f3-dc5d-491e-80d2-15c7ed5dd5c4/28c02623-42fa-4f5f-a984-
a02044bfc005' -H 'Accept: application/hal+json'
```

### Retrieve vVol bindings

GET `/protocols/san/vvol-bindings`

**Introduced In:** 9.10

Retrieves vVol bindings.

#### Related ONTAP commands

- `lun bind show`
- DOC /protocols/san/vvol-bindings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| secondary_id | string | query | False | Filter by secondary_id<br>• Introduced in: 9.13 |
| is_optimal | boolean | query | False | Filter by is_optimal |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| vvol.uuid | string | query | False | Filter by vvol.uuid |
| vvol.name | string | query | False | Filter by vvol.name |
| protocol_endpoint.uuid | string | query | False | Filter by protocol_endpoint.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol_endpoint.name | string | query | False | Filter by protocol_endpoint.name |
| count | integer | query | False | Filter by count |
| id | integer | query | False | Filter by id |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[vvol_binding] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "count": 1,
      "id": 1,
      "is_optimal": 1,
      "protocol_endpoint": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "secondary_id": "0000D20000010000h",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "vvol": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
          },
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

protocol_endpoint

The class `protocol_endpoint` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

svm

The SVM in which the vVol binding and its LUNs are located. Required in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vvol

The class `vvol` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

vvol_binding

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

Class `protocol_endpoint` and `vvol` LUNs support many-to-many vVol bindings. A LUN of one class can be bound to zero or more LUNs of the opposite class.

The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed. |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| is_optimal | boolean | Indicates if the class `protocol_endpoint` LUN and the class `vvol` LUN are on the same cluster node. |
| protocol_endpoint | protocol_endpoint | The class `protocol_endpoint` LUN in the vVol binding. Required in POST. |

| Name | Type | Description |
|---|---|---|
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |
| svm | svm | The SVM in which the vVol binding and its LUNs are located. Required in POST. |
| vvol | vvol | The class `vvol` LUN in the vVol binding. Required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a vVol binding**

POST `/protocols/san/vvol-bindings`

**Introduced In:** 9.10

Creates a vVol binding. The binding between any two specific LUNs is reference counted. When a binding is created that already exists, the binding count is incremented.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the vVol binding.
- `protocol_endpoint.uuid` or `protocol_endpoint.name` - Existing class `protocol_endpoint` LUN to bind to the specified class `vvol` LUN.
- `vvol.uuid` or `vvol.name` - Existing class `vvol` LUN to bind to the specified class `protocol_endpoint` LUN.

#### Related ONTAP commands

- `lun bind create`

#### Learn more

- DOC /protocols/san/vvol-bindings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| count | integer | The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed. |

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| is_optimal | boolean | Indicates if the class `protocol_endpoint` LUN and the class `vvol` LUN are on the same cluster node. |
| protocol_endpoint | protocol_endpoint | The class `protocol_endpoint` LUN in the vVol binding. Required in POST. |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |
| svm | svm | The SVM in which the vVol binding and its LUNs are located. Required in POST. |
| vvol | vvol | The class `vvol` LUN in the vVol binding. Required in POST. |

#### Example request

```
{
  "count": 1,
  "id": 1,
  "is_optimal": 1,
  "protocol_endpoint": {
    "name": "/vol/volume1/lun1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "secondary_id": "0000D20000010000h",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "vvol": {
    "name": "/vol/volume1/lun1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | The number of records in the response. |
| records | array[vvol_binding] | |

#### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "count": 1,
      "id": 1,
      "is_optimal": 1,
      "protocol_endpoint": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "secondary_id": "0000D20000010000h",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "vvol": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
```


#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621462 | The specified SVM does not exist or is not accessible to the caller. |
| 2621706 | Both the SVM UUID and SVM name were supplied, but don't refer to the same SVM. |

| Error Code | Description |
| --- | --- |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5374238 | A LUN in a snapshot was specified. |
| 5374323 | The LUN specified as the protocol endpoint LUN is not of class `protocol_endpoint`. |
| 5374325 | The LUN specified as the vVol LUN is not of class `vvol`. |
| 5374874 | The UUID and name supplied for the protocol endpoint of Vvol LUN do not refer to the same LUN. Use to the `target` property of the error object to differentiate between the protocol endpoint LUN and the vVol LUN. |
| 5374875 | The protocol endpoint or vVol LUN was not found or is not accessible to the caller. Use to the `target` property of the error object to differentiate between the protocol endpoint LUN and the vVol LUN. |
| 5374876 | The protocol endpoint or vVol LUN was not found in the SVM. Use to the `target` property of the error object to differentiate between the protocol endpoint LUN and the vVol LUN. |
| 5374924 | No protocol endpoint LUN was supplied. |
| 5374925 | No vVol LUN was supplied. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

protocol_endpoint

The class `protocol_endpoint` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
|---|---|---|
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

svm

The SVM in which the vVol binding and its LUNs are located. Required in POST.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vvol

The class `vvol` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

vvol_binding

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

Class `protocol_endpoint` and `vvol` LUNs support many-to-many vVol bindings. A LUN of one class can be bound to zero or more LUNs of the opposite class.

The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed. |

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| is_optimal | boolean | Indicates if the class `protocol_endpoint` LUN and the class `vvol` LUN are on the same cluster node. |
| protocol_endpoint | protocol_endpoint | The class `protocol_endpoint` LUN in the vVol binding. Required in POST. |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |
| svm | svm | The SVM in which the vVol binding and its LUNs are located. Required in POST. |
| vvol | vvol | The class `vvol` LUN in the vVol binding. Required in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a vVol binding

DELETE `/protocols/san/vvol-bindings/{protocol_endpoint.uuid}/{vvol.uuid}`

**Introduced In:** 9.10

Deletes a vVol binding. The binding between any two specific LUNs is reference counted. When a binding is deleted, the binding count is decremented, but the LUNs remain bound if the resultant reference count is greater than zero. When the binding count reaches zero, the binding is destroyed.

#### Related ONTAP commands

- `lun bind destroy`

#### Learn more

- DOC /protocols/san/vvol-bindings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol_endpoint.uuid | string | path | True | The unique identifier of the class `protocol_endpoint` LUN. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| vvol.uuid | string | path | True | The unique identifier of the class `vvol` LUN. |
| delete_all_references | boolean | query | False | Forces deletion of the binding regardless of the reference count value.<br>• Default value: |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The vVol binding was not found because the protocol endpoint or vVol LUN was not found. Use to the `target` property of the error object to differentiate between the protocol endpoint LUN and the vVol LUN. |
| 5374926 | The vVol binding was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a vVol binding

GET `/protocols/san/vvol-bindings/{protocol_endpoint.uuid}/{vvol.uuid}`

**Introduced In:** 9.10

Retrieves a vVol binding.

#### Related ONTAP commands

- `lun bind show`

#### Learn more

- DOC /protocols/san/vvol-bindings

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| protocol_endpoint.uuid | string | path | True | The unique identifier of the class `protocol_endpoint` LUN. |
| vvol.uuid | string | path | True | The unique identifier of the class `vvol` LUN. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| count | integer | The vVol binding between any two specific LUNs is reference counted. When a REST POST is executed for a vVol binding that already exists, the vVol binding reference count is incremented. When a REST DELETE is executed, the vVol binding reference count is decremented. Only when the vVol binding count reaches zero, or the query parameter `delete_all_references` is supplied, is the vVol binding destroyed. |

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| is_optimal | boolean | Indicates if the class `protocol_endpoint` LUN and the class `vvol` LUN are on the same cluster node. |
| protocol_endpoint | protocol_endpoint | The class `protocol_endpoint` LUN in the vVol binding. Required in POST. |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |
| svm | svm | The SVM in which the vVol binding and its LUNs are located. Required in POST. |
| vvol | vvol | The class `vvol` LUN in the vVol binding. Required in POST. |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "count": 1,
  "id": 1,
  "is_optimal": 1,
  "protocol_endpoint": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/lun1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "secondary_id": "0000D20000010000h",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "vvol": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/lun1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```


### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4 | The vVol binding was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

protocol_endpoint

The class `protocol_endpoint` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

svm

The SVM in which the vVol binding and its LUNs are located. Required in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

vvol

The class `vvol` LUN in the vVol binding. Required in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of a LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the LUN. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage LUNs

### Manage LUNs

#### Overview

A LUN is the logical representation of storage in a storage area network (SAN).

The LUN REST API allows you to create, update, delete, and discover LUNs.

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

#### Performance monitoring

Performance of a LUN can be monitored by observing the `metric.*` and `statistics.*` properties. These properties show the space utilization and performance of a LUN in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

#### Examples

### Creating a LUN

This example creates a 300 gigabyte, thin-provisioned LUN in SVM *svm1*, volume *vol1*, configured for use by *linux* initiators. The `return_records` query parameter is used to retrieve properties of the newly created LUN in the POST response.

```
# The API:
POST /api/storage/luns

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/luns?return_records=true' -H
'Accept: application/hal+json' -d '{ "svm": { "name": "svm1" }, "os_type":
"linux", "space": { "size": "300G" }, "name" : "/vol/vol1/lun1" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",
    "svm": {
      "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
        }
```

```
      }
    },
    "name": "/vol/vol1/lun1",
    "location": {
      "logical_unit": "lun1",
      "volume": {
        "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-
005056bbc17d"
          }
        }
      }
    },
    "class": "regular",
    "enabled": true,
    "os_type": "linux",
    "serial_number": "wf0Iq+N4uck3",
    "space": {
      "guarantee": {
        "requested": false,
        "reserved": false
      },
      "scsi_thin_provisioning_support_enabled": true,
      "size": 322163441664,
      "used": 0
    },
    "status": {
      "container_state": "online",
      "read_only": false,
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6"
      }
    }
  }
]
}
```


#### Updating a LUN

This example sets the `comment` property of a LUN.

```
# The API:
PATCH /api/storage/luns/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6' -H 'Accept: application/hal+json' -d '{ "comment": "Data for
the finance department." }'
```

#### Retrieving LUNs

This example retrieves summary information for all online LUNs in SVM *svm1*. The `svm.name` and `status.state` query parameters are used to find the desired LUNs.

```
# The API:
GET /api/storage/luns

# The call:
curl -X GET 'https://<mgmt-
ip>/api/storage/luns?svm.name=svm1&status.state=online' -H 'Accept:
application/hal+json'

# The response:
{
"records": [
  {
    "uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol1/lun1",
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6"
      }
    }
  },
  {
    "uuid": "c903a978-9bac-4ce9-8237-4a3ba8b13f08",
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol1/lun2",
```

```
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/luns/c903a978-9bac-4ce9-8237-4a3ba8b13f08"
      }
    }
  },
  {
    "uuid": "7faf0a9e-0a47-4876-8318-3638d5da16bf",
    "svm": {
      "name": "svm1"
    },
    "name": "/vol/vol2/lun3",
    "status": {
      "state": "online"
    },
    "_links": {
      "self": {
        "href": "/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/luns?svm.name=svm1&status.state=online"
  }
}
}
```


**Retrieving details for a specific LUN**

In this example, the `fields` query parameter is used to request all fields, including advanced fields, that would not otherwise be returned by default for the LUN.

```
# The API:
GET /api/storage/luns/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6?fields=**' -H 'Accept: application/hal+json'

# The response:
```

```
{
"uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",
"svm": {
  "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
    }
  }
},
"name": "/vol/vol1/lun1",
"location": {
  "logical_unit": "lun1",
  "volume": {
    "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-
005056bbc17d"
      }
    }
  }
},
"auto_delete": false,
"class": "vvol",
"comment": "Data for the finance department.",
"enabled": true,
"lun_maps": [
  {
    "logical_unit_number": 0,
    "igroup": {
      "uuid": "2b9d57e1-2a66-11e9-b682-005056bbc17d",
      "name": "ig1",
      "_links": {
        "self": {
          "href": "/api/protocols/san/igroups/2b9d57e1-2a66-11e9-b682-
005056bbc17d"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/lun-maps/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6/2b9d57e1-2a66-11e9-b682-005056bbc17d"
```

```
      }
    }
  }
],
"os_type": "linux",
"serial_number": "wf0Iq+N4uck3",
"space": {
  "guarantee": {
    "requested": false,
    "reserved": false
  },
  "scsi_thin_provisioning_support_enabled": true,
  "size": 322163441664,
  "used": 0
},
"vvol": {
  "is_bound": true,
  "bindings": [
    {
      "id": 4304512,
      "partner": {
        "uuid": "353c7262-be4b-4176-acf3-f1021faa8b64",
        "name": "/vol/vol1/pelun1",
        "_links": {
          "self": {
            "href": "/api/storage/luns/353c7262-be4b-4176-acf3-
f1021faa8b64"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/san/vvol-bindings/353c7262-be4b-4176-
acf3-f1021faa8b64/5a24ae5b-28af-47fb-b129-5adf6cfba0a6"
        }
      }
    }
  ]
},
 "metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 0,
```

```
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"status": {
  "container_state": "online",
  "mapped": true,
  "read_only": false,
  "state": "online"
},
"consistency_group": {
  "name": "vol1",
```

```
  "uuid": "6d657aaf-b57a-5396-82ea-c01329e46c79",
  "_links": {
    "self": {
      "href": "/api/application/consistency-groups/6d657aaf-b57a-5396-
82ea-c01329e46c79"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6?fields=**"
  }
}
}
```

**Deleting a LUN**

```
# The API:
DELETE /api/storage/luns/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/storage/luns/c903a978-9bac-4ce9-
8237-4a3ba8b13f08' -H 'Accept: application/hal+json'
```

---

#### LUN data

The LUN REST API also supports reading data from and writing data to a LUN via the REST API as multipart/form-data.

Reading data is performed using a GET request on the LUN endpoint. The request header must include `Accept: multipart/form-data`. When this header entry is provided, query parameters `data.offset` and `data.size` are required and used to specify the portion of the LUN's data to read; no other query parameters are allowed. Reads are limited to one megabyte (1MB) per request. Data is returned as `multipart/form-data` content with exactly one form entry containing the data. The form entry has content type `application/octet-stream`.

Writing data is performed using a PATCH request on the LUN endpoint. The request header must include `Content-Type: multipart/form-data`. When this header entry is provided, query parameter `data.offset` is required and used to specify the location within the LUN at which to write the data; no other query parameters are allowed. The request body must be `multipart/form-data` content with exactly one form entry containing the data to write. The content type entry of the form data is ignored and always treated as `application/octet-stream`. Writes are limited to one megabyte (1MB) per request.

#### Reading data from a LUN

```
# The API:
GET /api/storage/luns/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns/c903a978-9bac-4ce9-8237-
4a3ba8b13f08?data.offset=0&data.size=9' -H 'Accept: multipart/form-data'

# In the response header:
Content-Type: multipart/form-data; boundary="c6e9cf51ab354af0"

# The response body:
--c6e9cf51ab354af0
Content-Disposition: form-data;
Content-Type: application/octet-stream
data here
--c6e9cf51ab354af0--
```

#### Writing data to a LUN

This example writes the contents of a file to a LUN starting at offset 1024.

```
# The API:
PATCH /api/storage/luns/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/c903a978-9bac-4ce9-8237-
4a3ba8b13f08?data.offset=1024' -F "data=@file;type=application/octet-
stream"
```

---

#### Cloning LUNs

A clone of a LUN is an independent "copy" of the LUN that shares unchanged data blocks with the original. As blocks of the source and clone are modified, unique blocks are written for each. LUN clones can be created quickly and consume very little space initially. They can be created for the purpose of back-up, or to replicate data for multiple consumers.

Space reservations can be set for the LUN clone independent of the source LUN by setting the `space.guarantee.requested` property in a POST or PATCH request.

A LUN clone can be set to auto-delete by setting the `auto_delete` property. If the LUN's volume is configured for automatic deletion, LUNs that have auto-delete enabled are deleted when a volume is nearly full to reclaim a target amount of free space in the volume.

The value of property `space.scsi_thin_provisioning_support_enabled` is not propagated to the destination when a LUN is cloned as a new LUN; it is reset to false. The value of this property is maintained from the destination LUN when a LUN is overwritten as a clone.

#### Creating a new LUN clone

You create a new LUN clone as you create any LUN - a POST request to /storage/luns. Set `clone.source.uuid` or `clone.source.name` to identify the source LUN from which the clone is created. The LUN clone and its source must reside in the same volume.

The source LUN can reside in a snapshot, in which case the `clone.source.name` field must be used to identify it. Add `/.snapshot/<snapshot_name>` to the path after the volume name to identify the snapshot. For example `/vol/vol1/.snapshot/snap1/lun1`.

By default, new LUN clones do not inherit the QoS policy of the source LUN; a QoS policy should be set for the clone by setting the `qos_policy` property.

```
# The API:
POST /api/storage/luns

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/luns' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name":
"/vol/vol1/lun2clone1", "clone": { "source": { "name": "/vol/vol1/lun2" }
}, "qos_policy": { "name": "qos1" } }'
```

#### Over-writing an existing LUN's data as a clone of another

You can overwrite an existing LUN as a clone of another, using a PATCH request to /storage/luns/{uuid}. Set the `clone.source.uuid` or `clone.source.name` property to identify the source LUN from which the clone data is taken. The LUN clone and its source must reside in the same volume.

When used in a PATCH request, the patched LUN's data is overwritten as a clone of the source. The following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `vvol`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the updated LUN are also preserved.

```
# The API:
PATCH /api/storage/luns/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6' -H 'Accept: application/hal+json' -d '{ "clone": { "source":
{ "name": "/vol/vol1/lun2" } } }'
```

#### Converting an NVMe namespace into a LUN

An existing NVMe namespace can be converted in-place to a LUN with no modification to the data blocks. In other words, there is no additional copy created for the data blocks. There are certain requirements for converting an NVMe namespace to a LUN. For instance, the namespace should not be mapped to an NVMe subsystem. Additionally, the namespace should not have a block size other than 512 bytes.

The conversion process updates the metadata to the NVMe namespace, making it a LUN. The conversion is both time and space efficient. After conversion, the new LUN behaves as a regular LUN and may be mapped to an initiator group.

### Convert an NVMe namespace into a LUN

You convert an NVMe namespace into a LUN by calling a POST to /storage/luns. Set `convert.namespace.uuid` or `convert.namespace.name` to identify the source NVMe namespace which is to be converted in-place into a LUN.

```
# The API:
POST /api/storage/luns

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/luns' -H 'Accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "convert": {
"namespace": { "name": "/vol/vol1/namespace1" } } }'
```

---

#### Moving LUNs between volumes

You move a LUN between volumes by using a PATCH request to /storage/luns/{uuid}. Set the volume portion of the fully qualified LUN path `name` property, `path.volume.uuid`, or `path.volume.name` property to a different volume than the LUN's current volume. Moving a LUN between volumes is an asynchronous activity. A successful request returns a response of 200 synchronously, which indicates that the movement has been successfully queued. The LUN object can then be further polled with a GET request to `/storage/luns/{uuid}` to monitor the status of the movement.

The `movement` sub-object of the LUN object is populated while a LUN movement is in progress and for two minutes following completion of a movement.

### Starting a LUN movement

```
# The API:
PATCH /api/storage/luns/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/7faf0a9e-0a47-4876-8318-
3638d5da16bf' -H 'Accept: application/hal+json' -d '{ "name":
"/vol/vol1/lun3" }'
```

#### Checking on the status of the LUN movement

```
# The API:
GET /api/storage/luns/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns/7faf0a9e-0a47-4876-8318-
3638d5da16bf?fields=movement' -H 'Accept: application/hal+json'

# The response:
{
"uuid": "7faf0a9e-0a47-4876-8318-3638d5da16bf",
"name": "/vol/vol1/lun3",
"movement": {
  "paths": {
    "destination": "/vol/vol1/lun3",
    "source": "/vol/vol2/lun3"
  },
  "progress": {
    "elapsed": 1,
    "percent_complete": 0,
    "state": "preparing",
    "volume_snapshot_blocked": false
  }
},
"_links": {
  "self": {
    "href": "/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf"
  }
}
}
```

---

### Retrieve LUNs

GET `/storage/luns`

**Introduced In:** 9.6

Retrieves LUNs.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `attributes.*`
- `auto_delete`
- `copy.*`
- `lun_maps.*`
- `movement.*`
- `space.physical_used`
- `space.physical_used_by_snapshots`
- `space.efficiency_ratio`
- `space.snapshot.*`
- `statistics.*`
- `vvol.bindings.*`
- `metric.*`

#### Related ONTAP commands

- `lun bind show`
- `lun copy show`
- `lun mapping show`
- `lun move show`
- `lun show`
- `volume file clone show-autodelete`

#### Learn more

- DOC /storage/luns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| location.volume.name | string | query | False | Filter by location.volume.name |
| location.volume.uuid | string | query | False | Filter by location.volume.uuid |
| location.qtree.name | string | query | False | Filter by location.qtree.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| location.qtree.id | integer | query | False | Filter by location.qtree.id<br>• Max value: 4994<br>• Min value: 0 |
| location.logical_unit | string | query | False | Filter by location.logical_unit |
| location.node.name | string | query | False | Filter by location.node.name<br>• Introduced in: 9.10 |
| location.node.uuid | string | query | False | Filter by location.node.uuid<br>• Introduced in: 9.10 |
| enabled | boolean | query | False | Filter by enabled |
| name | string | query | False | Filter by name |
| statistics.status | string | query | False | Filter by statistics.status<br>• Introduced in: 9.7 |
| statistics.throughput_raw.other | integer | query | False | Filter by statistics.throughput_raw.other<br>• Introduced in: 9.7 |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write<br><br>• Introduced in: 9.7 |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total<br>• Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write<br>• Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read<br>• Introduced in: 9.7 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp<br>• Introduced in: 9.7 |
| class | string | query | False | Filter by class |
| space.physical_used | integer | query | False | Filter by space.physical_used<br>• Introduced in: 9.16 |
| space.size | integer | query | False | Filter by space.size<br>• Max value: 140737488355328<br>• Min value: 4096 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.scsi_thin_provisioning_support_enabled | boolean | query | False | Filter by space.scsi_thin_provisioning_support_enabled<br>• Introduced in: 9.10 |
| space.used | integer | query | False | Filter by space.used |
| space.efficiency_ratio | number | query | False | Filter by space.efficiency_ratio<br>• Introduced in: 9.16 |
| space.physical_used_by_snapshots | integer | query | False | Filter by space.physical_used_by_snapshots<br>• Introduced in: 9.16 |
| space.guarantee.reserved | boolean | query | False | Filter by space.guarantee.reserved |
| space.guarantee.requested | boolean | query | False | Filter by space.guarantee.requested |
| uuid | string | query | False | Filter by uuid |
| comment | string | query | False | Filter by comment<br>• maxLength: 254<br>• minLength: 0 |
| qos_policy.uuid | string | query | False | Filter by qos_policy.uuid |
| qos_policy.name | string | query | False | Filter by qos_policy.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.status | string | query | False | Filter by metric.status<br><br>• Introduced in: 9.7 |
| metric.duration | string | query | False | Filter by metric.duration<br><br>• Introduced in: 9.7 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other<br><br>• Introduced in: 9.7 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total<br><br>• Introduced in: 9.7 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write<br><br>• Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read<br><br>• Introduced in: 9.7 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other<br><br>• Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.iops.write | integer | query | False | Filter by metric.iops.write<br><br>• Introduced in: 9.7 |
| metric.iops.read | integer | query | False | Filter by metric.iops.read<br><br>• Introduced in: 9.7 |
| metric.timestamp | string | query | False | Filter by metric.timestamp<br><br>• Introduced in: 9.7 |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other<br><br>• Introduced in: 9.7 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total<br><br>• Introduced in: 9.7 |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write<br><br>• Introduced in: 9.7 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| attributes.value | string | query | False | Filter by attributes.value<br>• Introduced in: 9.10<br>• maxLength: 4091<br>• minLength: 1 |
| attributes.name | string | query | False | Filter by attributes.name<br>• Introduced in: 9.10<br>• maxLength: 4091<br>• minLength: 1 |
| lun_maps.igroup.name | string | query | False | Filter by lun_maps.igroup.name |
| lun_maps.igroup.initiators.name | string | query | False | Filter by lun_maps.igroup.initiators.name<br>• Introduced in: 9.16 |
| lun_maps.igroup.initiators.comment | string | query | False | Filter by lun_maps.igroup.initiators.comment<br>• Introduced in: 9.16<br>• maxLength: 254<br>• minLength: 0 |
| lun_maps.igroup.os_type | string | query | False | Filter by lun_maps.igroup.os_type<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun_maps.igroup.comment | string | query | False | Filter by lun_maps.igroup.comment<br>• Introduced in: 9.16<br>• maxLength: 254<br>• minLength: 0 |
| lun_maps.igroup.uuid | string | query | False | Filter by lun_maps.igroup.uuid |
| lun_maps.igroup.protocol | string | query | False | Filter by lun_maps.igroup.protocol<br>• Introduced in: 9.16 |
| lun_maps.igroup.igroups.uuid | string | query | False | Filter by lun_maps.igroup.igroups.uuid<br>• Introduced in: 9.16 |
| lun_maps.igroup.igroups.name | string | query | False | Filter by lun_maps.igroup.igroups.name<br>• Introduced in: 9.16<br>• maxLength: 96<br>• minLength: 1 |
| lun_maps.logical_unit_number | integer | query | False | Filter by lun_maps.logical_unit_number |
| status.container_state | string | query | False | Filter by status.container_state |
| status.mapped | boolean | query | False | Filter by status.mapped |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| status.read_only | boolean | query | False | Filter by status.read_only |
| status.state | string | query | False | Filter by status.state |
| vvol.is_bound | boolean | query | False | Filter by vvol.is_bound<br>• Introduced in: 9.10 |
| vvol.bindings.id | integer | query | False | Filter by vvol.bindings.id<br>• Introduced in: 9.10 |
| vvol.bindings.secondary_id | string | query | False | Filter by vvol.bindings.secondary_id<br>• Introduced in: 9.13 |
| vvol.bindings.partner.uuid | string | query | False | Filter by vvol.bindings.partner.uuid<br>• Introduced in: 9.10 |
| vvol.bindings.partner.name | string | query | False | Filter by vvol.bindings.partner.name<br>• Introduced in: 9.10 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| movement.paths.source | string | query | False | Filter by movement.paths.source |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| movement.paths.destination | string | query | False | Filter by movement.paths.destination |
| movement.progress.volume_snapshot_blocked | boolean | query | False | Filter by movement.progress.volume_snapshot_blocked |
| movement.progress.percent_complete | integer | query | False | Filter by movement.progress.percent_complete<br>• Max value: 100<br>• Min value: 0 |
| movement.progress.failure.arguments.message | string | query | False | Filter by movement.progress.failure.arguments.message |
| movement.progress.failure.arguments.code | string | query | False | Filter by movement.progress.failure.arguments.code |
| movement.progress.failure.code | string | query | False | Filter by movement.progress.failure.code |
| movement.progress.failure.message | string | query | False | Filter by movement.progress.failure.message |
| movement.progress.elapsed | integer | query | False | Filter by movement.progress.elapsed |
| movement.progress.state | string | query | False | Filter by movement.progress.state |
| movement.max_throughput | integer | query | False | Filter by movement.max_throughput |
| auto_delete | boolean | query | False | Filter by auto_delete |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| serial_number | string | query | False | Filter by serial_number<br>• maxLength: 12<br>• minLength: 12 |
| copy.destinations.name | string | query | False | Filter by copy.destinations.name<br>• Introduced in: 9.10 |
| copy.destinations.progress.percent_complete | integer | query | False | Filter by copy.destinations.progress.percent_complete<br>• Introduced in: 9.10<br>• Max value: 100<br>• Min value: 0 |
| copy.destinations.progress.volume_snapshot_blocked | boolean | query | False | Filter by copy.destinations.progress.volume_snapshot_blocked<br>• Introduced in: 9.10 |
| copy.destinations.progress.failure.arguments.message | string | query | False | Filter by copy.destinations.progress.failure.arguments.message<br>• Introduced in: 9.10 |
| copy.destinations.progress.failure.arguments.code | string | query | False | Filter by copy.destinations.progress.failure.arguments.code<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| copy.destinations.progress.failure.code | string | query | False | Filter by copy.destinations.progress.failure.code<br><br>• Introduced in: 9.10 |
| copy.destinations.progress.failure.message | string | query | False | Filter by copy.destinations.progress.failure.message<br><br>• Introduced in: 9.10 |
| copy.destinations.progress.elapsed | integer | query | False | Filter by copy.destinations.progress.elapsed<br><br>• Introduced in: 9.10 |
| copy.destinations.progress.state | string | query | False | Filter by copy.destinations.progress.state<br><br>• Introduced in: 9.10 |
| copy.destinations.peer.name | string | query | False | Filter by copy.destinations.peer.name<br><br>• Introduced in: 9.14 |
| copy.destinations.peer.uuid | string | query | False | Filter by copy.destinations.peer.uuid<br><br>• Introduced in: 9.14 |
| copy.destinations.max_throughput | integer | query | False | Filter by copy.destinations.max_throughput<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| copy.destinations.uuid | string | query | False | Filter by copy.destinations.uuid<br><br>• Introduced in: 9.10 |
| copy.source.max_throughput | integer | query | False | Filter by copy.source.max_throughput<br><br>• Introduced in: 9.10 |
| copy.source.uuid | string | query | False | Filter by copy.source.uuid<br><br>• Introduced in: 9.10 |
| copy.source.progress.failure.arguments.message | string | query | False | Filter by copy.source.progress.failure.arguments.message<br><br>• Introduced in: 9.10 |
| copy.source.progress.failure.arguments.code | string | query | False | Filter by copy.source.progress.failure.arguments.code<br><br>• Introduced in: 9.10 |
| copy.source.progress.failure.code | string | query | False | Filter by copy.source.progress.failure.code<br><br>• Introduced in: 9.10 |
| copy.source.progress.failure.message | string | query | False | Filter by copy.source.progress.failure.message<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| copy.source.progress.volume_snapshot_blocked | boolean | query | False | Filter by copy.source.progress.volume_snapshot_blocked<br>• Introduced in: 9.10 |
| copy.source.progress.percent_complete | integer | query | False | Filter by copy.source.progress.percent_complete<br>• Introduced in: 9.10<br>• Max value: 100<br>• Min value: 0 |
| copy.source.progress.state | string | query | False | Filter by copy.source.progress.state<br>• Introduced in: 9.10 |
| copy.source.progress.elapsed | integer | query | False | Filter by copy.source.progress.elapsed<br>• Introduced in: 9.10 |
| copy.source.peer.name | string | query | False | Filter by copy.source.peer.name<br>• Introduced in: 9.14 |
| copy.source.peer.uuid | string | query | False | Filter by copy.source.peer.uuid<br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| copy.source.name | string | query | False | Filter by copy.source.name<br>• Introduced in: 9.10 |
| consistency_group.name | string | query | False | Filter by consistency_group.name<br>• Introduced in: 9.10 |
| consistency_group.uuid | string | query | False | Filter by consistency_group.uuid<br>• Introduced in: 9.10 |
| serial_number_hex | string | query | False | Filter by serial_number_hex<br>• Introduced in: 9.17 |
| os_type | string | query | False | Filter by os_type |
| create_time | string | query | False | Filter by create_time<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Max value: 120<br>• Min value: 0<br>• Default value: 15 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[lun] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "attributes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "name1",
          "value": "value1"
        }
      ],
      "class": "string",
      "comment": "string",
      "consistency_group": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cg1",
        "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
      },
      "copy": {
        "destinations": [
          {
            "_links": {
              "self": {
```

```
                "href": "/api/resourcelink"
              }
            },
            "max_throughput": 0,
            "name": "/vol/vol1/lun1",
            "peer": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "peer1",
              "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
            },
            "progress": {
              "elapsed": 0,
              "failure": {
                "arguments": [
                  {
                    "code": "string",
                    "message": "string"
                  }
                ],
                "code": "4",
                "message": "entry doesn't exist"
              },
              "percent_complete": 0,
              "state": "string"
            },
            "uuid": "1bc327d5-4654-5284-a116-f182282240b4"
          }
        ],
        "source": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/vol2/lun1",
          "peer": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "peer1",
```

```
            "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
          },
          "progress": {
            "elapsed": 0,
            "failure": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
              ],
              "code": "4",
              "message": "entry doesn't exist"
            },
            "percent_complete": 0,
            "state": "string"
          },
          "uuid": "03c05019-40d9-3945-c767-dca4c3be5e90"
        }
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "location": {
        "logical_unit": "lun1",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "qtree": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "id": 1,
          "name": "qt1"
        },
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
```

```
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    "lun_maps": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "igroup": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "comment": "string",
          "igroups": [
            {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "igroup1",
              "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
            }
          ],
          "initiators": [
            {
              "comment": "my comment",
              "name": "iqn.1998-01.com.corp.iscsi:name1"
            }
          ],
          "name": "igroup1",
          "os_type": "string",
          "protocol": "string",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "metric": {
      "_links": {
        "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "movement": {
        "paths": {
          "destination": "/vol/vol1/lun1",
          "source": "/vol/vol2/lun2"
        },
        "progress": {
          "elapsed": 0,
          "failure": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          },
          "percent_complete": 0,
          "state": "string"
        }
      },
      "name": "/vol/volume1/qtree1/lun1",
      "os_type": "string",
      "qos_policy": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "qos1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "serial_number": "string",
      "serial_number_hex": "string",
      "space": {
        "efficiency_ratio": 2.5,
        "physical_used": 1073741824,
        "physical_used_by_snapshots": 1073741824,
        "size": 1073741824,
        "used": 0
      },
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "status": {
        "container_state": "string",
        "state": "online"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vvol": {
        "bindings": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "id": 1,
            "partner": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "/vol/vol1/lun1",
              "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
            },
            "secondary_id": "0000D20000010000h"
          }
        ]
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

attributes

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

Optional in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The attribute name. |
| value | string | The attribute value. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`,

`qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

consistency_group

The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported. A LUN's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

namespace

The source namespace for convert operation. This can be specified using property `convert.namespace.uuid` or `convert.namespace.name`. If both properties are supplied, they must refer to the same NVMe namespace.

Valid in POST. A convert request from NVMe namespace to LUN cannot be combined with setting any other LUN properties. All other properties of the converted LUN come from the source NVMe namespace.

convert

This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is desired.

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

This is only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

Error information provided if the asynchronous LUN copy operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| failure | error | Error information provided if the asynchronous LUN copy operation fails. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy. |

| Name | Type | Description |
| --- | --- | --- |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

destinations

A LUN copy operation in which the containing LUN is the source of the copy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. See property `copy.source.max_throughput` for further details. |
| name | string | The fully qualified path of the LUN copy destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>This is only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy destination. |

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
|---|---|---|
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| failure | error | Error information provided if the asynchronous LUN copy operation fails. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy.<br><br>Valid in PATCH when an LUN copy is active. Set to *paused* to pause a LUN copy. Set to *replicating* to resume a paused LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

source

The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.

Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN copy, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>Valid only in a POST that begins a LUN copy or a PATCH when a LUN copy is already in process.<br><br>• format: int64<br>• Introduced in: 9.10<br>• x-nullable: true |
| name | string | The fully qualified path of the LUN copy source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

| Name | Type | Description |
| --- | --- | --- |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy source.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

copy

This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.

Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.

While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.

There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| destinations | array[destinations] | An array of destination LUNs of LUN copy operations in which the containing LUN is the source of the copy. |
| source | source | The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.<br><br>Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN. |

node

The cluster node that hosts the LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the LUN is located. Valid in POST and PATCH.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the LUN within the ONTAP cluster.
LUNs support rename and move between volumes. Valid in POST and PATCH.

| Name | Type | Description |
|---|---|---|
| logical_unit | string | The base name component of the LUN. Valid in POST and PATCH.<br><br>If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.<br><br>A PATCH that modifies the base name of the LUN is considered a rename operation. |
| node | node | The cluster node that hosts the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | The qtree in which the LUN is optionally located. Valid in POST and PATCH.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.<br><br>A PATCH that modifies the qtree of the LUN is considered a rename operation. |
| volume | volume | The volume in which the LUN is located. Valid in POST and PATCH.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

igroups

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group to which the LUN is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| igroup | igroup | The initiator group to which the LUN is mapped. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br><br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

paths

The fully qualified LUN path names involved in the LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| destination | string | The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| source | string | The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

error

Error information provided if the asynchronous LUN movement operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN movement, in seconds. |
| failure | error | Error information provided if the asynchronous LUN movement operation fails. |
| percent_complete | integer | The percentage completed of the LUN movement. |
| state | string | The state of the LUN movement.<br><br>Valid in PATCH when an LUN movement is active. Set to *paused* to pause a LUN movement. Set to *replicating* to resume a paused LUN movement. |

| Name | Type | Description |
| --- | --- | --- |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN movement. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN movement. |

movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.<br><br>• format: int64<br>• Introduced in: 9.6<br>• x-nullable: true |
| paths | paths | The fully qualified LUN path names involved in the LUN movement. |
| progress | progress | Properties related to the progress of an active or recently completed LUN movement. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

qos_policy

The QoS policy for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. This property is mutually exclusive with the LUN granular

`qos_policy`. If no `qos_policy` is provided at LUN or volume granularity, a volume granular policy is set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST.

object_stores

tiering

The tiering placement and policy definitions for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
|---|---|---|
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST. |

qos_policy

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| uuid | string | The unique identifier of the QoS policy. Valid in POST and PATCH. |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.<br>This property is caller settable as described above. |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_ratio | number | The storage efficiency ratio of the LUN without snapshots. (Logical Used / Used)<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |
| physical_used | integer | The number of bytes consumed on the disk by the LUN, excluding snapshots.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the LUN's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| scsi_thin_provisioning_support_enabled | boolean | To leverage the benefits of SCSI thin provisioning, it must be supported by your host. SCSI thin provisioning uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard. Only hosts that support this standard can use SCSI thin provisioning in ONTAP.<br><br>When you disable SCSI thin provisioning support in ONTAP, you turn off the following SCSI thin provisioning features:<br><br>• Unmapping and reporting space usage for space reclamation<br>• Reporting resource exhaustion errors<br><br>The value of this property is not propagated to the destination when a LUN is cloned as a new LUN or copied; it is reset to false. The value of this property is maintained from the destination LUN when a LUN is overwritten as a clone.<br><br>Valid in POST and PATCH.<br><br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not decreased using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the LUN.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available. |
| mapped | boolean | Reports if the LUN is mapped to one or more initiator groups.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the LUN allows only read access. |
| state | string | The state of the LUN. Normal states for a LUN are *online* and *offline*. Other states indicate errors. |

svm

The SVM in which the LUN is located.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

partner

The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the partner LUN. |
| uuid | string | The unique identifier of the partner LUN. |

bindings

A vVol binding with which the LUN is associated.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| partner | partner | The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN. |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |

vvol

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.

There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| bindings | array[bindings] | Bindings between the LUN, which must be of class `protocol_endpoint` or `vvol`, and LUNs of the opposite class.<br><br>A class `vvol` LUN must be bound to a class `protocol_endpoint` LUN in order to be accessed. Class `protocol_endpoint` and `vvol` LUNs allow many-to-many bindings. A LUN of one class is allowed to be bound to zero or more LUNs of the opposite class. The binding between any two specific LUNs is reference counted. When a binding is created that already exists, the binding count is incremented. When a binding is deleted, the binding count is decremented, but the LUNs remain bound if the resultant reference count is greater than zero. When the binding count reaches zero, the binding is destroyed.<br><br>The bindings array contains LUNs of the opposite class of the containing LUN object.<br><br>There is an added computational cost to retrieving property values for `vvol.bindings`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| is_bound | boolean | Reports if the LUN is part of a VMware virtual volume (vVol) bind relationship. This is `true` if the LUN is of class `protocol_endpoint` or `vvol` and has one or more bindings to a LUN of the opposite class. This is false if the LUN is of class `regular` or unbound. |

lun

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| attributes | array[attributes] | An array of name/value pairs optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.<br><br>Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.<br><br>Valid in POST except when creating a LUN clone. A cloned can already have attributes from its source. You can add, modify, and delete the attributes of a LUN clone in separate requests after creation of the LUN.<br><br>Attributes can be added/modified/removed for an existing LUN using the /api/storage/luns/{lun.uuid}/attributes endpoint. For further information, see DOC /storage/luns/{lun.uuid}/attributes .<br><br>There is an added computational cost to retrieving property values for `attributes`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br><br>• Introduced in: 9.10<br>• readCreate: 1 |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| class | string | The class of LUN.<br><br>Optional in POST. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, check the `status.state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |
| lun_maps | array[lun_maps] | The LUN maps with which the LUN is associated.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| metric | metric | Performance numbers, such as IOPS latency and throughput. |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | status | Status information about the LUN. |
| svm | svm | The SVM in which the LUN is located. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a LUN**

POST `/storage/luns`

**Introduced In:** 9.6

Creates a LUN.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the LUN.
- `name`, `location.volume.name` or `location.volume.uuid` - Existing volume in which to create the LUN.
- `name` or `location.logical_unit` - Base name of the LUN.
- `os_type` - Operating system from which the LUN will be accessed. Required when creating a non-clone LUN and disallowed when creating a clone of an existing LUN. A clone's `os_type` is taken from the source LUN.
- `space.size` - Size of the LUN. Required when creating a non-clone LUN and disallowed when creating a clone of an existing LUN. A clone's size is taken from the source LUN.

#### Recommended optional properties

- `qos_policy.name` or `qos_policy.uuid` - Existing traditional or adaptive QoS policy to be applied to the LUN. All LUNs should be managed by a QoS policy at the volume or LUN level.

#### Default property values

If not specified in POST, the follow default property values are assigned.

- `auto_delete` - *false*

#### Related ONTAP commands

- `lun create`
- `lun convert-from-namespace`
- `lun copy start`
- `volume file clone autodelete`
- `volume file clone create`

POST is asynchronous when creating a new LUN. It is synchronous when converting a namespace to a LUN via the `convert` property.

#### Learn more

- DOC /storage/luns

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| attributes | array[attributes] | An array of name/value pairs optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.<br><br>Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.<br><br>Valid in POST except when creating a LUN clone. A cloned can already have attributes from its source. You can add, modify, and delete the attributes of a LUN clone in separate requests after creation of the LUN.<br><br>Attributes can be added/modified/removed for an existing LUN using the /api/storage/luns/{lun.uuid}/attributes endpoint. For further information, see DOC /storage/luns/{lun.uuid}/attributes .<br><br>There is an added computational cost to retrieving property values for `attributes`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br><br>• Introduced in: 9.10<br>• readCreate: 1 |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| class | string | The class of LUN.<br><br>Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |
| convert | convert | This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is desired. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |
| lun_maps | array[lun_maps] | The LUN maps with which the LUN is associated.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| status | status | Status information about the LUN. |
| svm | svm | The SVM in which the LUN is located. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

**Example request**

```
{
  "attributes": [
    {
      "name": "name1",
      "value": "value1"
    }
  ],
  "class": "string",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "consistency_group": {
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "convert": {
    "namespace": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "copy": {
    "destinations": [
      {
        "max_throughput": 0,
        "name": "/vol/vol1/lun1",
        "peer": {
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "progress": {
          "elapsed": 0,
          "percent_complete": 0,
          "state": "string"
        },
        "uuid": "1bc327d5-4654-5284-a116-f182282240b4"
      }
    ],
    "source": {
      "name": "/vol/vol2/lun1",
```

```
      "peer": {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      },
      "uuid": "03c05019-40d9-3945-c767-dca4c3be5e90"
    }
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "enabled": true,
  "location": {
    "logical_unit": "lun1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": {
      "id": 1,
      "name": "qt1"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "lun_maps": [
    {
      "igroup": {
        "comment": "string",
        "igroups": [
          {
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "initiators": [
          {
            "comment": "my comment",
            "name": "iqn.1998-01.com.corp.iscsi:name1"
          }
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
```

```
  ],
  "movement": {
    "paths": {
      "destination": "/vol/vol1/lun1",
      "source": "/vol/vol2/lun2"
    }
  },
  "name": "/vol/volume1/qtree1/lun1",
  "os_type": "string",
  "provisioning_options": {
    "exclude_aggregates": [
      {
        "name": "aggr1"
      }
    ],
    "qos_policy": {
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "snapshot_policy": {
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "storage_service": {
      "name": "string"
    },
    "tiering": {
      "control": "string",
      "object_stores": [
        {
          "name": "string"
        }
      ],
      "policy": "string"
    }
  },
  "qos_policy": {
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "serial_number_hex": "string",
  "space": {
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
```

```
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vvol": {
    "bindings": [
      {
        "id": 1,
        "partner": {
          "name": "/vol/vol1/lun1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "secondary_id": "0000D20000010000h"
      }
    ]
  }
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

**Example response**


```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917927 | The specified volume was not found. |
| 918235 | A volume with the specified UUID was not found. |
| 918236 | The specified `location.volume.uuid` and `location.volume.name` do not refer to the same volume. |
| 1260121 | Cloning a LUN to a volume different than the source volume is not supported. |
| 1260136 | The specified destination for a clone operation already exists as a LUN, namespace, or file. |
| 2621462 | The specified SVM does not exist. |
| 2621706 | The specified `svm.uuid` and `svm.name` do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied. |
| 5242927 | The specified qtree was not found. |
| 5242950 | The specified `location.qtree.id` and `location.qtree.name` do not refer to the same qtree. |
| 5374121 | A LUN name can only contain characters A-Z, a-z, 0-9, "-", ".", "_", "{" and "}". |
| 5374123 | A negative size was provided for the LUN. |
| 5374124 | The specified size is too small for the LUN. |
| 5374125 | The specified size is too large for the LUN. |

| Error Code | Description |
| --- | --- |
| 5374127 | The specified LUN name is invalid. |
| 5374129 | LUNs cannot be created on a load sharing mirror volume. |
| 5374130 | An invalid size value was provided. |
| 5374237 | LUNs cannot be created on an SVM root volume. |
| 5374238 | LUNs cannot be created in snapshots. |
| 5374241 | A size value with invalid units was provided. |
| 5374242 | A LUN or NVMe namespace already exists at the specified path. |
| 5374352 | An invalid name was provided for the LUN. |
| 5374614 | The specified storage availability zone does not exist. |
| 5374707 | Creating a LUN in the specific volume is not allowed because the volume is reserved for an application. |
| 5374858 | The volume specified by `name` is not the same as that specified by `location.volume`. |
| 5374859 | No volume was specified for the LUN. |
| 5374860 | The qtree specified by `name` is not the same as that specified by `location.qtree`. |
| 5374861 | The LUN base name specified by `name` is not the same as that specified by `location.logical_unit`. |
| 5374862 | No LUN path base name was provided for the LUN. |
| 5374863 | An error occurred after successfully creating the LUN. Some properties were not set. |
| 5374874 | The specified `clone.source.uuid` and `clone.source.name` do not refer to the same LUN. |
| 5374875 | The specified `clone.source` was not found. |
| 5374876 | The specified `clone.source` was not found. |
| 5374883 | The property cannot be specified when creating a LUN clone. The `target` property of the error object identifies the property. |
| 5374884 | A property that is required when creating a new LUN that is not a LUN clone or LUN copy was not supplied. The `target` property of the error object identifies the property. |
| 5374886 | An error occurred after successfully creating the LUN preventing the retrieval of its properties. |
| 5374899 | The `clone.source.uuid` property is not supported when specifying a source LUN from a snapshot. |

| Error Code | Description |
| --- | --- |
| 5374928 | An incomplete attribute name/value pair was supplied. |
| 5374929 | The combined sizes of an attribute name and value are too large. |
| 5374932 | A name for an attribute was duplicated. |
| 5374942 | The property cannot be specified at the same time when creating a LUN as a clone. The `target` property of the error object identifies the other property given with clone. |
| 5374943 | The property cannot be specified at the same time when creating a LUN as a copy. The `target` property of the error object identifies the other property given with copy. |
| 5374944 | The property cannot be specified when converting an NVMe namespace into a LUN. The `target` property of the error object identifies the property. |
| 5375054 | The source LUN is required when requesting an SVM copy operation. |
| 5375059 | An unsuitable QoS policy was specified. |
| 5375061 | The specified `location.storage_availability_zone.uuid` and `location.storage_availability_zone.name` do not refer to the same storage availability zone. |
| 5376461 | The specified LUN name is invalid. |
| 5376462 | The specified LUN name is too long. |
| 5376463 | The snapshot portion of the specified LUN name is too long. |
| 5376469 | The property cannot be set during the LUN create operation on this platform. |
| 5376515 | The specified storage availability zone is not configured for provisioning on the specified SVM. |
| 5440509 | No suitable storage can be found for the specified requirements. |
| 5440688 | A LUN cannot be created with the same name as an existing LUN. |
| 5702832 | A LUN or namespace with the same name is already being created. LUN and namespace names must be unique within an SVM. |
| 7018877 | Maximum combined total (50) of file and LUN copy and move operations reached. When one or more of the operations has completed, try the command again. |

| Error Code | Description |
| --- | --- |
| 13565952 | The LUN clone request failed. |
| 26345672 | An SVM peer relationship was not found with the provided 'copy.source.peer.uuid'. |
| 26345673 | An SVM peer relationship between the SVM provided in 'copy.source.peer.name' and the SVM provided in either 'svm.name' or 'svm.uuid' was not found. |
| 26345674 | The SVM peer relationship provided in 'copy.source.peer.uuid' does not match a peer relationship between the local SVM provided in either 'svm.uuid' or 'svm.name' and the peer SVM provided in 'copy.source.peer.name'. |
| 26345675 | The destination SVM provided via 'svm.name' or 'svm.uuid' does not match the local SVM for the provided SVM peer relationship UUID in 'copy.source.peer.uuid'. |
| 72089755 | NVMe namespace with a block size of 4096 bytes cannot be converted to a LUN. |
| 72089756 | Namespace is currently mapped to subsystem. |
| 72089757 | NVMe namespace in a snapshot cannot be converted to a LUN. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

attributes

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

Optional in POST.

| Name | Type | Description |
|---|---|---|
| name | string | The attribute name. |
| value | string | The attribute value. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the clone source LUN.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

| Name | Type | Description |
| --- | --- | --- |
| source | source | The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.<br><br>Valid in POST to create a new LUN as a clone of the source.<br><br>Valid in PATCH to overwrite an existing LUN's data as a clone of another. |

consistency_group

The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported. A LUN's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

namespace

The source namespace for convert operation. This can be specified using property `convert.namespace.uuid` or `convert.namespace.name`. If both properties are supplied, they must refer to the same NVMe namespace.

Valid in POST. A convert request from NVMe namespace to LUN cannot be combined with setting any other LUN properties. All other properties of the converted LUN come from the source NVMe namespace.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the source NVMe namespace. Valid in POST.<br>An NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>NVMe namespace names are paths of the form "/vol/<volume>[/<qtree>]/<NVMe namespace>" where the qtree name is optional. |
| uuid | string | The unique identifier of the source NVMe namespace. Valid in POST. |

convert

This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is desired.

| Name | Type | Description |
|---|---|---|
| namespace | namespace | The source namespace for convert operation. This can be specified using property `convert.namespace.uuid` or `convert.namespace.name`. If both properties are supplied, they must refer to the same NVMe namespace.<br><br>Valid in POST. A convert request from NVMe namespace to LUN cannot be combined with setting any other LUN properties. All other properties of the converted LUN come from the source NVMe namespace. |

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

This is only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

Error information provided if the asynchronous LUN copy operation fails.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
|---|---|---|
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

destinations

A LUN copy operation in which the containing LUN is the source of the copy.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. See property `copy.source.max_throughput` for further details. |
| name | string | The fully qualified path of the LUN copy destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>This is only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy destination. |

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
|---|---|---|
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

source

The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.

Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN copy, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>Valid only in a POST that begins a LUN copy or a PATCH when a LUN copy is already in process.<br><br>• format: int64<br>• Introduced in: 9.10<br>• x-nullable: true |
| name | string | The fully qualified path of the LUN copy source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

| Name | Type | Description |
| --- | --- | --- |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM. |
| uuid | string | The unique identifier of the LUN copy source.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

copy

This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.

Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.

While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.

There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| destinations | array[destinations] | An array of destination LUNs of LUN copy operations in which the containing LUN is the source of the copy. |

| Name | Type | Description |
| --- | --- | --- |
| source | source | The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.<br><br>Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN. |

node

The cluster node that hosts the LUN.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the LUN is located. Valid in POST and PATCH.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the LUN within the ONTAP cluster.
LUNs support rename and move between volumes. Valid in POST and PATCH.

| Name | Type | Description |
| --- | --- | --- |
| logical_unit | string | The base name component of the LUN. Valid in POST and PATCH.<br><br>If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.<br><br>A PATCH that modifies the base name of the LUN is considered a rename operation. |
| node | node | The cluster node that hosts the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | The qtree in which the LUN is optionally located. Valid in POST and PATCH.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.<br><br>A PATCH that modifies the qtree of the LUN is considered a rename operation. |
| volume | volume | The volume in which the LUN is located. Valid in POST and PATCH.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation. |

self_link

igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group to which the LUN is mapped.

| Name | Type | Description |
|---|---|---|
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |

| Name | Type | Description |
| --- | --- | --- |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| igroup | igroup | The initiator group to which the LUN is mapped. |

| Name | Type | Description |
| --- | --- | --- |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

paths

The fully qualified LUN path names involved in the LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| destination | string | The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| source | string | The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

error

Error information provided if the asynchronous LUN movement operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN movement, in seconds. |
| percent_complete | integer | The percentage completed of the LUN movement. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN movement. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN movement. |

movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
|---|---|---|
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.<br><br>• format: int64<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| paths | paths | The fully qualified LUN path names involved in the LUN movement. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

qos_policy

The QoS policy for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. This property is mutually exclusive with the LUN granular `qos_policy`. If no `qos_policy` is provided at LUN or volume granularity, a volume granular policy is set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Storage service name. If not specified, the default value is the most performant for the platform. |

object_stores

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
| --- | --- | --- |
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| auto | boolean | If the volume specified in the request does not exist, automatically provision one of appropriate size. If the volume does exist, resize it to accommodate the new LUN.<br><br>This property is only supported on Unified ONTAP.<br><br>The following behavior is different from a traditional POST request:<br><br>• The operation is asynchronous.<br>• The `qos_policy` property is applied to the provisioned volume instead of the LUN. A default QoS policy is applied to the volume if one is not provided.<br>• The `provisioning_options.count` property is supported, provisioning *count* LUNs on the volume using the specified properties.<br>• The `lun_maps` property is supported. If the specified initiator group does not exist, it is created. The LUN is mapped to this initiator group. If an initiator group is provisioned in this way, it is deleted after it is no longer mapped to any LUNs.<br>• The `clone`, `copy`, and `convert` properties are not supported.<br>• When performing `records` based operations, specifying this property in the query applies to the entire operation. Specifying it for an individual record within the request applies to only that record.<br>• Many other `provisioning_options` properties are supported to control the placement of the LUN and the properties of the volume containing the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| count | integer | The number of LUNs to provision with these properties. Only POST requests based on `space.size` are supported. When provided, the name is considered a prefix, and a suffix of the form _<*N*> is generated where N is the next available numeric index, starting with 1. |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |
| qos_policy | qos_policy | The QoS policy for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. This property is mutually exclusive with the LUN granular `qos_policy`. If no `qos_policy` is provided at LUN or volume granularity, a volume granular policy is set based on the `storage_service.name`, which defaults to the most performant service available. |
| snapshot_policy | snapshot_policy_reference | This is a reference to the snapshot policy. |
| storage_service | storage_service | Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST. |
| tiering | tiering | The tiering placement and policy definitions for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. |

| Name | Type | Description |
| --- | --- | --- |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning the volume to host the LUN. Only mirrored aggregates are used if this parameter is set to *true* and only unmirrored aggregates are used if this parameter is set to *false*. The default value is *true* for a MetroCluster configuration and is *false* for a non-MetroCluster configuration. |

qos_policy

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| uuid | string | The unique identifier of the QoS policy. Valid in POST and PATCH. |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.<br>This property is caller settable as described above. |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_ratio | number | The storage efficiency ratio of the LUN without snapshots. (Logical Used / Used)<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| physical_used | integer | The number of bytes consumed on the disk by the LUN, excluding snapshots.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the LUN's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| scsi_thin_provisioning_support_enabled | boolean | To leverage the benefits of SCSI thin provisioning, it must be supported by your host. SCSI thin provisioning uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard. Only hosts that support this standard can use SCSI thin provisioning in ONTAP.<br><br>When you disable SCSI thin provisioning support in ONTAP, you turn off the following SCSI thin provisioning features:<br><br>• Unmapping and reporting space usage for space reclamation<br>• Reporting resource exhaustion errors<br><br>The value of this property is not propagated to the destination when a LUN is cloned as a new LUN or copied; it is reset to false. The value of this property is maintained from the destination LUN when a LUN is overwritten as a clone.<br><br>Valid in POST and PATCH.<br><br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not decreased using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the LUN.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available. |
| mapped | boolean | Reports if the LUN is mapped to one or more initiator groups.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the LUN allows only read access. |
| state | string | The state of the LUN. Normal states for a LUN are *online* and *offline*. Other states indicate errors. |

svm

The SVM in which the LUN is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

partner

The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the partner LUN. |
| uuid | string | The unique identifier of the partner LUN. |

bindings

A vVol binding with which the LUN is associated.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| partner | partner | The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN. |

| Name | Type | Description |
| --- | --- | --- |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |

vvol

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.

There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| bindings | array[bindings] | Bindings between the LUN, which must be of class `protocol_endpoint` or `vvol`, and LUNs of the opposite class.<br><br>A class `vvol` LUN must be bound to a class `protocol_endpoint` LUN in order to be accessed. Class `protocol_endpoint` and `vvol` LUNs allow many-to-many bindings. A LUN of one class is allowed to be bound to zero or more LUNs of the opposite class. The binding between any two specific LUNs is reference counted. When a binding is created that already exists, the binding count is incremented. When a binding is deleted, the binding count is decremented, but the LUNs remain bound if the resultant reference count is greater than zero. When the binding count reaches zero, the binding is destroyed.<br><br>The bindings array contains LUNs of the opposite class of the containing LUN object.<br><br>There is an added computational cost to retrieving property values for `vvol.bindings`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| is_bound | boolean | Reports if the LUN is part of a VMware virtual volume (vVol) bind relationship. This is `true` if the LUN is of class `protocol_endpoint` or `vvol` and has one or more bindings to a LUN of the opposite class. This is false if the LUN is of class `regular` or unbound. |

lun

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
| --- | --- | --- |
| attributes | array[attributes] | An array of name/value pairs optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.<br><br>Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.<br><br>Valid in POST except when creating a LUN clone. A cloned can already have attributes from its source. You can add, modify, and delete the attributes of a LUN clone in separate requests after creation of the LUN.<br><br>Attributes can be added/modified/removed for an existing LUN using the /api/storage/luns/{lun.uuid}/attributes endpoint. For further information, see DOC /storage/luns/{lun.uuid}/attributes .<br><br>There is an added computational cost to retrieving property values for `attributes`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br><br>• Introduced in: 9.10<br>• readCreate: 1 |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| class | string | The class of LUN.<br><br>Optional in POST. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |

| Name | Type | Description |
| --- | --- | --- |
| convert | convert | This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is desired. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |
| lun_maps | array[lun_maps] | The LUN maps with which the LUN is associated.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| provisioning_options | provisioning_options | Options that are applied to the operation. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| status | status | Status information about the LUN. |
| svm | svm | The SVM in which the LUN is located. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete a LUN

DELETE `/storage/luns/{uuid}`

**Introduced In:** 9.6

Deletes a LUN.

#### Related ONTAP commands

- `lun copy cancel`
- `lun delete`

#### Learn more

- DOC /storage/luns

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | The unique identifier of the LUN to retrieve. |
| allow_delete_while_mapped | boolean | query | False | Allows deletion of a mapped LUN.<br><br>A mapped LUN might be in use. Deleting a mapped LUN also deletes the LUN map and makes the data no longer available. This might cause a disruption in the availability of data.<br><br>**This parameter should be used with caution.**<br><br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1254197 | The LUN is mapped and cannot be deleted without specifying the `allow_delete_while_mapped` query parameter. |
| 5374705 | Deleting the LUN is not allowed because it is part of an application. |
| 5374865 | The LUN's aggregate is offline. The aggregate must be online to modify or remove the LUN. |
| 5374866 | The LUN's volume is offline. The volume must be online to modify or remove the LUN. |
| 5374875 | The specified LUN was not found. |
| 5374876 | The specified LUN was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve LUN properties or data

GET `/storage/luns/{uuid}`

**Introduced In:** 9.6

Retrieves a LUN's properties or a LUN's data.

LUN data read requests are distinguished by the header entry `Accept: multipart/form-data`. When this header entry is provided, query parameters `data.offset` and `data.size` are required and used to specify the portion of the LUN's data to read; no other query parameters are allowed. Reads are limited to one megabyte (1MB) per request. Data is returned as `multipart/form-data` content with exactly one form entry containing the data. The form entry has content type `application/octet-stream`.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `attributes.*`
- `auto_delete`
- `copy.*`
- `lun_maps.*`

- `movement.*`
- `space.physical_used`
- `space.physical_used_by_snapshots`
- `space.efficiency_ratio`
- `space.snapshot.*`
- `statistics.*`
- `vvol.bindings.*`
- `metric.*`

#### Related ONTAP commands

- `lun bind show`
- `lun copy show`
- `lun mapping show`
- `lun move show`
- `lun show`
- `volume file clone show-autodelete`

#### Learn more

- DOC /storage/luns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the LUN to retrieve. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| data.offset | integer | query | False | The offset, in bytes, at which to begin reading LUN data.<br><br>LUN data read requests are distinguished by the header entry `Accept: multipart/form-data`. When this header entry is provided, query parameters `data.offset` and `data.size` are required and used to specify the portion of the LUN's data to read; no other query parameters are allowed. Reads are limited to one megabyte (1MB) per request. Data is returned as `multipart/form-data` content with exactly one form entry containing the data. The form entry has content type `application/octet-stream`.<br><br>• format: int64<br>• Min value: 0<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| data.size | integer | query | False | The size, in bytes, of LUN data to read.<br><br>LUN data read requests are distinguished by the header entry `Accept: multipart/form-data`. When this header entry is provided, query parameters `data.offset` and `data.size` are required and used to specify the portion of the LUN's data to read; no other query parameters are allowed. Reads are limited to one megabyte (1MB) per request. Data is returned as `multipart/form-data` content with exactly one form entry containing the data. The form entry has content type `application/octet-stream`.<br><br>• format: int64<br>• Min value: 1<br>• Introduced in: 9.11<br>• Max value: 1048576 |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| attributes | array[attributes] | An array of name/value pairs optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.<br><br>Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.<br><br>Valid in POST except when creating a LUN clone. A cloned can already have attributes from its source. You can add, modify, and delete the attributes of a LUN clone in separate requests after creation of the LUN.<br><br>Attributes can be added/modified/removed for an existing LUN using the /api/storage/luns/{lun.uuid}/attributes endpoint. For further information, see DOC /storage/luns/{lun.uuid}/attributes .<br><br>There is an added computational cost to retrieving property values for `attributes`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br><br>• Introduced in: 9.10<br>• readCreate: 1 |

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| class | string | The class of LUN.<br><br>Optional in POST. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, check the `status.state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |
| lun_maps | array[lun_maps] | The LUN maps with which the LUN is associated.<br><br>There is an added computational cost to retrieving property values for `lun_maps`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.<br>These properties are supported for GET and POST. During POST, it requires the `provisioning_options.auto` property to be set to true.<br>See the `provisioning_options.auto` property for full details. |
| metric | metric | Performance numbers, such as IOPS latency and throughput. |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| os_type | string | The operating system type of the LUN.<br><br>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | status | Status information about the LUN. |
| svm | svm | The SVM in which the LUN is located. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "attributes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "name1",
      "value": "value1"
    }
  ],
  "class": "string",
  "comment": "string",
  "consistency_group": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "copy": {
    "destinations": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": 0,
        "name": "/vol/vol1/lun1",
        "peer": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "progress": {
          "elapsed": 0,
          "failure": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          },
          "percent_complete": 0,
          "state": "string"
        },
        "uuid": "1bc327d5-4654-5284-a116-f182282240b4"
      }
    ],
    "source": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/vol2/lun1",
      "peer": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      },
      "progress": {
        "elapsed": 0,
        "failure": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
```

```
          "code": "4",
          "message": "entry doesn't exist"
        },
        "percent_complete": 0,
        "state": "string"
      },
      "uuid": "03c05019-40d9-3945-c767-dca4c3be5e90"
    }
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "location": {
    "logical_unit": "lun1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 1,
      "name": "qt1"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "lun_maps": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "comment": "string",
        "igroups": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "igroup1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "initiators": [
          {
            "comment": "my comment",
            "name": "iqn.1998-01.com.corp.iscsi:name1"
          }
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
```

```
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "movement": {
    "paths": {
      "destination": "/vol/vol1/lun1",
      "source": "/vol/vol2/lun2"
    },
    "progress": {
      "elapsed": 0,
      "failure": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "percent_complete": 0,
      "state": "string"
    }
  },
  "name": "/vol/volume1/qtree1/lun1",
  "os_type": "string",
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "serial_number_hex": "string",
  "space": {
```

```
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
    "size": 1073741824,
    "used": 0
  },
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vvol": {
    "bindings": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "id": 1,
        "partner": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/vol1/lun1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "secondary_id": "0000D20000010000h"
      }
    ]
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN was not found. |
| 5374876 | The specified LUN was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

attributes

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

Optional in POST.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The attribute name. |
| value | string | The attribute value. |

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

consistency_group

The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported. A LUN's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

namespace

The source namespace for convert operation. This can be specified using property `convert.namespace.uuid` or `convert.namespace.name`. If both properties are supplied, they must refer to the same NVMe namespace.

Valid in POST. A convert request from NVMe namespace to LUN cannot be combined with setting any other LUN properties. All other properties of the converted LUN come from the source NVMe namespace.

convert

This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is desired.

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

This is only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

Error information provided if the asynchronous LUN copy operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| failure | error | Error information provided if the asynchronous LUN copy operation fails. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

destinations

A LUN copy operation in which the containing LUN is the source of the copy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. See property `copy.source.max_throughput` for further details. |
| name | string | The fully qualified path of the LUN copy destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>This is only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy destination. |

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| failure | error | Error information provided if the asynchronous LUN copy operation fails. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy.<br><br>Valid in PATCH when an LUN copy is active. Set to *paused* to pause a LUN copy. Set to *replicating* to resume a paused LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

source

The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.

Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN copy, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>Valid only in a POST that begins a LUN copy or a PATCH when a LUN copy is already in process.<br><br>• format: int64<br>• Introduced in: 9.10<br>• x-nullable: true |
| name | string | The fully qualified path of the LUN copy source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

| Name | Type | Description |
| --- | --- | --- |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy source.<br><br>Set this property in POST to specify the source for a LUN copy operation. |

copy

This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.

Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.

While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.

There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| destinations | array[destinations] | An array of destination LUNs of LUN copy operations in which the containing LUN is the source of the copy. |
| source | source | The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.<br><br>Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN. |

node

The cluster node that hosts the LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the LUN is located. Valid in POST and PATCH.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the LUN within the ONTAP cluster.
LUNs support rename and move between volumes. Valid in POST and PATCH.

| Name | Type | Description |
|---|---|---|
| logical_unit | string | The base name component of the LUN. Valid in POST and PATCH.<br><br>If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.<br><br>A PATCH that modifies the base name of the LUN is considered a rename operation. |
| node | node | The cluster node that hosts the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| qtree | qtree | The qtree in which the LUN is optionally located. Valid in POST and PATCH.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.<br><br>A PATCH that modifies the qtree of the LUN is considered a rename operation. |
| volume | volume | The volume in which the LUN is located. Valid in POST and PATCH.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation. |

self_link

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

igroups

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group to which the LUN is mapped.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment available for use by the administrator. Valid in POST and PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| igroups | array[igroups] | The existing initiator groups that are members of the group. Optional in POST.<br><br>This property is mutually exclusive with the *initiators* property during POST.<br><br>This array contains only the direct children of the initiator group. If the member initiator groups have further nested initiator groups, those are reported in the `igroups` property of the child initiator group.<br><br>Zero or more nested initiator groups can be supplied when the initiator group is created. The initiator group will act as if it contains the aggregation of all initiators in any nested initiator groups.<br><br>After creation, nested initiator groups can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/igroups` endpoint. See DELETE /protocols/san/igroups/{igroup.uuid}/igroups/{uuid} for more details. |
| initiators | array[initiators] | The initiators that are members of the group. |
| name | string | The name of the initiator group. |
| os_type | string | The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. |

| Name | Type | Description |
| --- | --- | --- |
| protocol | string | The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to *mixed*.<br><br>The protocol of an initiator group cannot be changed after creation of the group. |
| uuid | string | The unique identifier of the initiator group. |

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| igroup | igroup | The initiator group to which the LUN is mapped. |
| logical_unit_number | integer | The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through the Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value. This property is not supported when the *provisioning_options.count* property is 2 or more.<br><br>• Introduced in: 9.6<br>• readCreate: 1<br>• x-nullable: true |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

paths

The fully qualified LUN path names involved in the LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| destination | string | The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| source | string | The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

error

Error information provided if the asynchronous LUN movement operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN movement, in seconds. |
| failure | error | Error information provided if the asynchronous LUN movement operation fails. |
| percent_complete | integer | The percentage completed of the LUN movement. |
| state | string | The state of the LUN movement.<br><br>Valid in PATCH when an LUN movement is active. Set to *paused* to pause a LUN movement. Set to *replicating* to resume a paused LUN movement. |

| Name | Type | Description |
| --- | --- | --- |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN movement. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN movement. |

movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.<br><br>• format: int64<br>• Introduced in: 9.6<br>• x-nullable: true |
| paths | paths | The fully qualified LUN path names involved in the LUN movement. |
| progress | progress | Properties related to the progress of an active or recently completed LUN movement. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

qos_policy

The QoS policy for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. This property is mutually exclusive with the LUN granular

`qos_policy`. If no `qos_policy` is provided at LUN or volume granularity, a volume granular policy is set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST.

object_stores

tiering

The tiering placement and policy definitions for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
|---|---|---|
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold.<br><br>FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br><br>all ‐ Allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br><br>auto ‐ Allows tiering of both snapshot and active file system user data to the cloud store<br><br>none ‐ Volume blocks are not be tiered to the cloud store.<br><br>snapshot_only ‐ Allows tiering of only the volume snapshots not associated with the active file system.<br><br>The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy it is 31 days. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
| --- | --- | --- |
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| storage_service | storage_service | Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST. |

qos_policy

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| uuid | string | The unique identifier of the QoS policy. Valid in POST and PATCH. |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.<br>This property is caller settable as described above. |

| Name | Type | Description |
| --- | --- | --- |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_ratio | number | The storage efficiency ratio of the LUN without snapshots. (Logical Used / Used)<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |
| physical_used | integer | The number of bytes consumed on the disk by the LUN, excluding snapshots.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the LUN's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| scsi_thin_provisioning_support_enabled | boolean | To leverage the benefits of SCSI thin provisioning, it must be supported by your host. SCSI thin provisioning uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard. Only hosts that support this standard can use SCSI thin provisioning in ONTAP.<br><br>When you disable SCSI thin provisioning support in ONTAP, you turn off the following SCSI thin provisioning features:<br><br>• Unmapping and reporting space usage for space reclamation<br>• Reporting resource exhaustion errors<br><br>The value of this property is not propagated to the destination when a LUN is cloned as a new LUN or copied; it is reset to false. The value of this property is maintained from the destination LUN when a LUN is overwritten as a clone.<br><br>Valid in POST and PATCH.<br><br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not decreased using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the LUN.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available. |
| mapped | boolean | Reports if the LUN is mapped to one or more initiator groups.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the LUN allows only read access. |
| state | string | The state of the LUN. Normal states for a LUN are *online* and *offline*. Other states indicate errors. |

svm

The SVM in which the LUN is located.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

partner

The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the partner LUN. |
| uuid | string | The unique identifier of the partner LUN. |

bindings

A vVol binding with which the LUN is associated.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| partner | partner | The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN. |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |

vvol

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.

There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| bindings | array[bindings] | Bindings between the LUN, which must be of class `protocol_endpoint` or `vvol`, and LUNs of the opposite class.<br><br>A class `vvol` LUN must be bound to a class `protocol_endpoint` LUN in order to be accessed. Class `protocol_endpoint` and `vvol` LUNs allow many-to-many bindings. A LUN of one class is allowed to be bound to zero or more LUNs of the opposite class. The binding between any two specific LUNs is reference counted. When a binding is created that already exists, the binding count is incremented. When a binding is deleted, the binding count is decremented, but the LUNs remain bound if the resultant reference count is greater than zero. When the binding count reaches zero, the binding is destroyed.<br><br>The bindings array contains LUNs of the opposite class of the containing LUN object.<br><br>There is an added computational cost to retrieving property values for `vvol.bindings`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| is_bound | boolean | Reports if the LUN is part of a VMware virtual volume (vVol) bind relationship. This is `true` if the LUN is of class `protocol_endpoint` or `vvol` and has one or more bindings to a LUN of the opposite class. This is false if the LUN is of class `regular` or unbound. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an existing LUN

PATCH `/storage/luns/{uuid}`

**Introduced In:** 9.6

Updates an existing LUN in one of several ways:

- Updates the properties of a LUN.
- Writes data to a LUN. LUN data write requests are distinguished by the header entry `Content-Type: multipart/form-data`. When this header entry is provided, query parameter `data.offset` is required and used to specify the location within the LUN at which to write the data; no other query parameters are allowed. The request body must be `multipart/form-data` content with exactly one form entry containing the data to write. The content type entry of the form data is ignored and always treated as `application/octet-stream`. Writes are limited to one megabyte (1MB) per request.
- Overwrites the contents of a LUN as a clone of another.
- Begins the movement of a LUN between volumes. PATCH can also pause and resume the movement of a LUN between volumes that is already in active.

#### Related ONTAP commands

- `lun copy modify`
- `lun copy pause`

- `lun copy resume`
- `lun modify`
- `lun move-in-volume`
- `lun move modify`
- `lun move pause`
- `lun move resume`
- `lun move start`
- `lun resize`
- `volume file clone autodelete`

#### Learn more

- DOC /storage/luns

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | The unique identifier of the LUN to retrieve. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| data.offset | integer | query | False | The offset, in bytes, at which to begin writing LUN data.<br><br>LUN data write requests are distinguished by the header entry `Content-Type: multipart/form-data`. When this header entry is provided, query parameter `data.offset` is required and used to specify the location within the LUN at which to write the data; no other query parameters are allowed. The request body must be `multipart/form-data` content with exactly one form entry containing the data to write. The content type entry of the form data is ignored and always treated as `application/octet-stream`. Writes are limited to one megabyte (1MB) per request.<br><br>• format: int64<br>• Min value: 0<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Min value: 0<br>• Default value: 0<br>• Max value: 120 |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, check the `status.state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br><br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| status | status | Status information about the LUN. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

**Example request**

```
{
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "consistency_group": {
    "name": "cg1",
    "uuid": "4abc2317-4332-9d37-93a0-20bd29c22df0"
  },
  "copy": {
    "destinations": [
      {
        "max_throughput": 0,
        "name": "/vol/vol1/lun1",
        "peer": {
          "name": "peer1",
          "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
        },
        "progress": {
          "elapsed": 0,
          "percent_complete": 0,
          "state": "string"
        },
        "uuid": "1bc327d5-4654-5284-a116-f182282240b4"
      }
    ],
    "source": {
      "peer": {
        "name": "peer1",
        "uuid": "4204cf77-4c82-9bdb-5644-b5a841c097a9"
      },
      "progress": {
        "elapsed": 0,
        "percent_complete": 0,
        "state": "string"
      }
    }
  },
  "create_time": "2018-06-04 19:00:00 +0000",
  "location": {
    "logical_unit": "lun1",
```

```
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": {
      "id": 1,
      "name": "qt1"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "movement": {
    "paths": {
      "destination": "/vol/vol1/lun1",
      "source": "/vol/vol2/lun2"
    },
    "progress": {
      "elapsed": 0,
      "percent_complete": 0,
      "state": "string"
    }
  },
  "name": "/vol/volume1/qtree1/lun1",
  "qos_policy": {
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "serial_number_hex": "string",
  "space": {
    "efficiency_ratio": 2.5,
    "physical_used": 1073741824,
    "physical_used_by_snapshots": 1073741824,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vvol": {
    "bindings": [
      {
```

```
        "id": 1,
        "partner": {
          "name": "/vol/vol1/lun1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "secondary_id": "0000D20000010000h"
      }
    ]
  }
}
```


**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917927 | The specified volume was not found. |
| 918235 | A volume with the specified UUID was not found. |
| 918236 | The specified `location.volume.uuid` and `location.volume.name` do not refer to the same volume. |
| 1254196 | A LUN already exists at the specified path. |
| 1254226 | Cannot patch a LUN's size to its current size. |
| 5242927 | The specified qtree was not found. |
| 5242950 | The specified `location.qtree.id` and `location.qtree.name` do not refer to the same qtree. |
| 5374124 | The specified LUN size is too small. |
| 5374125 | The specified LUN size is too large. |

| Error Code | Description |
| --- | --- |
| 5374127 | The specified LUN name is invalid. |
| 5374130 | An invalid size value was provided. |
| 5374241 | A size value with invalid units was provided. |
| 5374480 | Modifying the LUN is not allowed because it is in a foreign LUN import relationship. |
| 5374858 | The volume specified by `name` is not the same as that specified by `location.volume`. |
| 5374860 | The qtree specified by `name` is not the same as that specified by `location.qtree`. |
| 5374861 | The LUN base name specified by `name` is not the same as that specified by `location.logical_unit`. |
| 5374864 | An error occurred after successfully overwriting data for the LUN as a clone. Some properties were not modified. |
| 5374865 | The LUN's aggregate is offline. The aggregate must be online to modify or remove the LUN. |
| 5374866 | The LUN's volume is offline. The volume must be online to modify or remove the LUN. |
| 5374874 | The specified `clone.source.uuid` and `clone.source.name` do not refer to the same LUN. |
| 5374875 | The specified LUN was not found. This can apply to `clone.source` or the target LUN. The `target` property of the error object identifies the property. |
| 5374876 | The specified LUN was not found. This can apply to `clone.source` or the target LUN. The `target` property of the error object identifies the property. |
| 5374885 | An error occurred after successfully modifying some of the properties of the LUN. Some properties were not modified. |
| 5374889 | An invalid value was specified for `movement.progress.state`. Active LUN movement operations can be PATCHed to only *paused* or *replicating*. |
| 5374892 | An attempt was made to reduce the size of a LUN. |
| 5374904 | The destination volume is not online. |
| 5375059 | An unsuitable QoS policy was specified. |
| 5376461 | The specified LUN name is invalid. |
| 5376462 | The specified LUN name is too long. |

| Error Code | Description |
| --- | --- |
| 5376463 | The snapshot portion of the specified LUN name is too long. |
| 5376466 | An attempt was made to rename a LUN to a snapshot name. |
| 5376467 | An attempt was made to rename a primary LUN to a secondary name. |
| 5376468 | An attempt was made to rename a LUN to a reserved name. |
| 5376470 | The property cannot be set during the LUN modify operation on this platform. |
| 7018877 | Maximum combined total (50) of file and LUN copy and move operations reached. When one or more of the operations has completed, try the command again. |
| 7018919 | A copy or move job exists with the same destination LUN. |
| 8454243 | The specified QoS policy does not exist while modifying a LUN. |
| 13565952 | The LUN clone request failed. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

attributes

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

Optional in POST.

source

The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the clone source LUN.<br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>" where the qtree name is optional.<br><br>Valid in POST and PATCH. |
| uuid | string | The unique identifier of the clone source LUN.<br>Valid in POST and PATCH. |

clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`,

`qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

| Name | Type | Description |
|---|---|---|
| source | source | The source LUN for a LUN clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same LUN.<br><br>Valid in POST to create a new LUN as a clone of the source.<br><br>Valid in PATCH to overwrite an existing LUN's data as a clone of another. |

consistency_group

The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported. A LUN's consistency group is the consistency group of its containing volume.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the consistency group. |
| uuid | string | The unique identifier of the consistency group. |

namespace

The source namespace for convert operation. This can be specified using property `convert.namespace.uuid` or `convert.namespace.name`. If both properties are supplied, they must refer to the same NVMe namespace.

Valid in POST. A convert request from NVMe namespace to LUN cannot be combined with setting any other LUN properties. All other properties of the converted LUN come from the source NVMe namespace.

convert

This sub-object is used in POST to convert a valid in-place NVMe namespace to a LUN. Setting a property in this sub-object indicates that a conversion from the specified NVMe namespace to LUN is

desired.

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

This is only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error

Error information provided if the asynchronous LUN copy operation fails.

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
|---|---|---|
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

destinations

A LUN copy operation in which the containing LUN is the source of the copy.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. See property `copy.source.max_throughput` for further details. |
| name | string | The fully qualified path of the LUN copy destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>This is only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |
| uuid | string | The unique identifier of the LUN copy destination. |

peer

The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.

Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM.

| Name | Type | Description |
|---|---|---|
| name | string | The local name of the peer SVM. This name is unique among all local and peer SVMs. |
| uuid | string | The unique identifier of the SVM peer relationship. This is the UUID of the relationship, not the UUID of the peer SVM itself. |

progress

Properties related to the progress of an active or recently completed LUN copy.

| Name | Type | Description |
|---|---|---|
| elapsed | integer | The amount of time that has elapsed since the start of the LUN copy, in seconds. |
| percent_complete | integer | The percentage completed of the LUN copy. |
| state | string | The state of the LUN copy.<br><br>Valid in PATCH when an LUN copy is active. Set to *paused* to pause a LUN copy. Set to *replicating* to resume a paused LUN copy. |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN copy. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN copy. |

source

The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.

Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting

any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN.

| Name | Type | Description |
|---|---|---|
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN copy. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN copy, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>Valid only in a POST that begins a LUN copy or a PATCH when a LUN copy is already in process.<br><br>• format: int64<br>• Introduced in: 9.10<br>• x-nullable: true |
| peer | peer | The SVM peer relationship object for an inter-SVM LUN copy operation. The peer SVM in the relationship is the source SVM and the local SVM is the destination SVM.<br><br>Set this in POST to specify the source SVM for an inter-SVM LUN copy. Only populated by GET when the LUN copy is inter-SVM. |
| progress | progress | Properties related to the progress of an active or recently completed LUN copy. |

copy

This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.

Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.

While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.

There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
|---|---|---|
| destinations | array[destinations] | An array of destination LUNs of LUN copy operations in which the containing LUN is the source of the copy. |
| source | source | The source LUN of a LUN copy operation in which the containing LUN is the destination of the copy.<br><br>Valid in POST except when creating a LUN clone. A LUN copy request cannot be combined with setting any other LUN properties except the destination location. All other properties of the destination LUN come from the source LUN. |

node

The cluster node that hosts the LUN.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

| Name | Type | Description |
|---|---|---|
| id | integer | The identifier for the qtree, unique within the qtree's volume. |
| name | string | The name of the qtree. |

volume

The volume in which the LUN is located. Valid in POST and PATCH.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

location

The location of the LUN within the ONTAP cluster.
LUNs support rename and move between volumes. Valid in POST and PATCH.

| Name | Type | Description |
| --- | --- | --- |
| logical_unit | string | The base name component of the LUN. Valid in POST and PATCH.<br><br>If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.<br><br>A PATCH that modifies the base name of the LUN is considered a rename operation. |
| node | node | The cluster node that hosts the LUN. |
| qtree | qtree | The qtree in which the LUN is optionally located. Valid in POST and PATCH.<br><br>If properties `name` and `location.qtree.name` and/or `location.qtree.id` are specified in the same request, they must refer to the same qtree.<br><br>A PATCH that modifies the qtree of the LUN is considered a rename operation. |
| volume | volume | The volume in which the LUN is located. Valid in POST and PATCH.<br><br>If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.<br><br>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation. |

self_link

igroups

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the initiator group. |
| uuid | string | The unique identifier of the initiator group. |

initiators

The initiators that are members of the initiator group.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment available for use by the administrator. |
| name | string | Name of initiator that is a member of the initiator group. |

igroup

The initiator group to which the LUN is mapped.

lun_maps

A LUN map is an association between a LUN and an initiator group.

When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

paths

The fully qualified LUN path names involved in the LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| destination | string | The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |
| source | string | The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the optional qtree name, and base name of the LUN. |

error

Error information provided if the asynchronous LUN movement operation fails.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

progress

Properties related to the progress of an active or recently completed LUN movement.

| Name | Type | Description |
| --- | --- | --- |
| elapsed | integer | The amount of time that has elapsed since the start of the LUN movement, in seconds. |
| percent_complete | integer | The percentage completed of the LUN movement. |
| state | string | The state of the LUN movement.<br><br>Valid in PATCH when an LUN movement is active. Set to *paused* to pause a LUN movement. Set to *replicating* to resume a paused LUN movement. |

| Name | Type | Description |
| --- | --- | --- |
| volume_snapshot_blocked | boolean | This property reports if volume snapshots are blocked by the LUN movement. This property can be polled to identify when volume snapshots can be resumed after beginning a LUN movement. |

movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | integer | The maximum data throughput, in bytes per second, that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.<br><br>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.<br><br>• format: int64<br>• Introduced in: 9.6<br>• x-nullable: true |
| paths | paths | The fully qualified LUN path names involved in the LUN movement. |
| progress | progress | Properties related to the progress of an active or recently completed LUN movement. |

exclude_aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | The aggregate name. |

qos_policy

The QoS policy for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume. This property is mutually exclusive with the LUN granular

`qos_policy`. If no `qos_policy` is provided at LUN or volume granularity, a volume granular policy is set based on the `storage_service.name`, which defaults to the most performant service available.

| Name | Type | Description |
|---|---|---|
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

snapshot_policy_reference

This is a reference to the snapshot policy.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

storage_service

Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST.

object_stores

| Name | Type | Description |
|---|---|---|
| name | string | The name of the object store to use. Used for placement. |

tiering

The tiering placement and policy definitions for the volume provisioned to host the LUN. This property is only supported when the request provisions a new volume.

| Name | Type | Description |
|---|---|---|
| control | string | Storage tiering placement rules for the object. |
| object_stores | array[object_stores] | Object stores to use. Used for placement. |

provisioning_options

Options that are applied to the operation.

| Name | Type | Description |
|---|---|---|
| exclude_aggregates | array[exclude_aggregates] | A list of aggregates to exclude when determining the placement of the volume. |
| storage_service | storage_service | Determines the placement of the LUN based on the value specified. This property is only supported for regular and vvol LUNs. Valid in POST. |

qos_policy

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| uuid | string | The unique identifier of the QoS policy. Valid in POST and PATCH. |

guarantee

Properties that request and report the space guarantee for the LUN.

| Name | Type | Description |
| --- | --- | --- |
| requested | boolean | The requested space reservation policy for the LUN. If *true*, a space reservation is requested for the LUN; if *false*, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.<br>This property is caller settable as described above. |
| reserved | boolean | Reports if the LUN is space guaranteed.<br><br>If *true*, a space guarantee is requested and the containing volume and aggregate support the request. If *false*, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request. |

space

The storage space related properties of the LUN.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_ratio | number | The storage efficiency ratio of the LUN without snapshots. (Logical Used / Used)<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| guarantee | guarantee | Properties that request and report the space guarantee for the LUN. |

| Name | Type | Description |
| --- | --- | --- |
| physical_used | integer | The number of bytes consumed on the disk by the LUN, excluding snapshots.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |
| physical_used_by_snapshots | integer | The number of bytes consumed on the disk by the LUN's snapshots.<br>This property has been replaced by `space.snapshot.used`.<br>This property is not available on the LUN object in the REST API and is not reported for GET requests. See the containing volume object for this information. |

| Name | Type | Description |
| --- | --- | --- |
| scsi_thin_provisioning_support_enabled | boolean | To leverage the benefits of SCSI thin provisioning, it must be supported by your host. SCSI thin provisioning uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard. Only hosts that support this standard can use SCSI thin provisioning in ONTAP.<br><br>When you disable SCSI thin provisioning support in ONTAP, you turn off the following SCSI thin provisioning features:<br><br>• Unmapping and reporting space usage for space reclamation<br>• Reporting resource exhaustion errors<br><br>The value of this property is not propagated to the destination when a LUN is cloned as a new LUN or copied; it is reset to false. The value of this property is maintained from the destination LUN when a LUN is overwritten as a clone.<br><br>Valid in POST and PATCH.<br><br>• Default value: 1<br>• Introduced in: 9.10<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | The total provisioned size of the LUN. The LUN size can be increased but not decreased using the REST interface.<br><br>The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes, in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• example: 1073741824<br>• format: int64<br>• Max value: 140737488355328<br>• Min value: 4096<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The amount of space consumed by the main data stream of the LUN.<br><br>This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.<br><br>For more information, see *Size properties* in the *docs* section of the ONTAP REST API documentation.<br><br>• format: int64<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

status

Status information about the LUN.

| Name | Type | Description |
| --- | --- | --- |
| container_state | string | The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available. |
| mapped | boolean | Reports if the LUN is mapped to one or more initiator groups.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| read_only | boolean | Reports if the LUN allows only read access. |
| state | string | The state of the LUN. Normal states for a LUN are *online* and *offline*. Other states indicate errors. |

svm

The SVM in which the LUN is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

partner

The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the partner LUN. |
| uuid | string | The unique identifier of the partner LUN. |

bindings

A vVol binding with which the LUN is associated.

| Name | Type | Description |
| --- | --- | --- |
| id | integer | The ONTAP internal identifier assigned to the vVol binding. The bind identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>This property was included in early releases of the REST API for vVols and is maintained for backward compatibility. See the `secondary_id` property, which replaces `id`.<br><br>• example: 1<br>• readOnly: 1<br>• x-ntap-deprecated: 9.13.1<br>• Introduced in: 9.10<br>• x-nullable: true |
| partner | partner | The LUN partner that this LUN is bound to. If this LUN is a `vvol` class LUN, the partner is a `protocol_endpoint` class LUN. |

| Name | Type | Description |
| --- | --- | --- |
| secondary_id | string | The identifier assigned to the vVol binding, known as the secondary LUN ID. The identifier is unique amongst all class `vvol` LUNs bound to the same class `protocol_endpoint` LUN.<br><br>The format for a secondary LUN ID is 16 hexadecimal digits (zero-filled) followed by a lower case "h". |

vvol

A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.

See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.

There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

| Name | Type | Description |
| --- | --- | --- |
| bindings | array[bindings] | Bindings between the LUN, which must be of class `protocol_endpoint` or `vvol`, and LUNs of the opposite class.<br><br>A class `vvol` LUN must be bound to a class `protocol_endpoint` LUN in order to be accessed. Class `protocol_endpoint` and `vvol` LUNs allow many-to-many bindings. A LUN of one class is allowed to be bound to zero or more LUNs of the opposite class. The binding between any two specific LUNs is reference counted. When a binding is created that already exists, the binding count is incremented. When a binding is deleted, the binding count is decremented, but the LUNs remain bound if the resultant reference count is greater than zero. When the binding count reaches zero, the binding is destroyed.<br><br>The bindings array contains LUNs of the opposite class of the containing LUN object.<br><br>There is an added computational cost to retrieving property values for `vvol.bindings`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| is_bound | boolean | Reports if the LUN is part of a VMware virtual volume (vVol) bind relationship. This is `true` if the LUN is of class `protocol_endpoint` or `vvol` and has one or more bindings to a LUN of the opposite class. This is false if the LUN is of class `regular` or unbound. |

lun

A LUN is the logical representation of storage in a storage area network (SAN).

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, moved to a different volume and copied. LUNs support the assignment of a QoS policy for performance management or a QoS policy can be assigned to a volume containing one or more LUNs.

| Name | Type | Description |
|---|---|---|
| auto_delete | boolean | This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.<br><br>When set to *true*, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.<br><br>This property is optional in POST and PATCH. The default value for a new LUN is *false*.<br><br>There is an added computational cost to retrieving this property's value. It is not populated for a GET request unless it is explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| clone | clone | This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, `space.guarantee.requested` and `space.scsi_thin_provisioning_support_enabled`.<br><br>When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.<br><br>Persistent reservations for the patched LUN are also preserved. |
| comment | string | A configurable comment available for use by the administrator. Valid in POST and PATCH. |
| consistency_group | consistency_group | The LUN's consistency group. This property is populated for LUNs that are members of a consistency group. If the LUN is a member of a child consistency group, the parent consistency group is reported.<br>A LUN's consistency group is the consistency group of its containing volume. |

| Name | Type | Description |
| --- | --- | --- |
| copy | copy | This sub-object applies to LUN copy operations. A LUN can be copied with a POST request that supplies `copy.source` properties.<br><br>Copying a LUN is an asynchronous activity begun by a POST request that specifies the source of the copy in the `copy.source` properties. The data for the LUN is then asynchronously copied from the source to the destination. The time required to complete the copy depends on the size of the LUN and the load on the cluster. The `copy` sub-object is populated while a LUN copy is in progress and for two (2) minutes following completion of a copy.<br><br>While LUNs are being copied, the status of the LUN copy operations can be obtained using a GET of the source or destination LUN that requests the `copy` properties. If the LUN is the source LUN for one or more copy operations, the `copy.destinations` array is populated in GET. If the containing LUN is the destination LUN for a copy operation, the `copy.source` sub-object is populated in GET. The LUN copy operation can be further modified using a PATCH on the properties on the `copy.source` sub-object of the copy destination LUN.<br><br>There is an added computational cost to retrieving property values for `copy`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |
| create_time | string | The time the LUN was created. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, check the `status.state` property to determine if the LUN is administratively disabled (*offline*) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to *true* or brought administratively offline by setting the `enabled` property to *false*. Upon creation, a LUN is enabled by default. Valid in PATCH. |
| location | location | The location of the LUN within the ONTAP cluster.<br>LUNs support rename and move between volumes. Valid in POST and PATCH.<br><br>• Introduced in: 9.6 |

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.<br><br>Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.<br><br>While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.<br><br>There is an added computational cost to retrieving property values for `movement`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the LUN. Valid in POST and PATCH.<br><br>A LUN is located within a volume. Optionally, it can be located within a qtree in a volume.<br><br>LUN names are paths of the form "/vol/<volume>[/<qtree>]/<lun>" where the qtree name is optional.<br><br>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.<br><br>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation. |
| qos_policy | qos_policy | The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set the property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH. |
| serial_number | string | The LUN serial number. The serial number is generated by ONTAP when the LUN is created.<br><br>• maxLength: 12<br>• minLength: 12<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| serial_number_hex | string | The LUN serial number encoded in hexadecimal format. The serial number is generated by ONTAP when the LUN is created.<br>• readOnly: 1<br>• Introduced in: 9.17<br>• x-nullable: true |
| space | space | The storage space related properties of the LUN. |
| status | status | Status information about the LUN. |
| uuid | string | The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.<br>• example: 1cd8a442-86d1-11e0-ae1c-123478563412<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| vvol | vvol | A VMware virtual volume (vVol) binding is an association between a LUN of class `protocol_endpoint` and a LUN of class `vvol`. Class `protocol_endpoint` LUNs are mapped to igroups and granted access using the same configuration as class `regular` LUNs. When a class `vvol` LUN is bound to a mapped class `protocol_endpoint` LUN, VMware can access the class `vvol` LUN through the class `protocol_endpoint` LUN mapping.<br><br>See DELETE /protocols/san/vvol-bindings to learn more about deleting vVol bindings.<br><br>There is an added computational cost to retrieving property values for `vvol`. They are not populated for a GET request unless explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage LUN attributes

## Manage LUN attributes

### Overview

LUN attributes are caller-defined name/value pairs optionally stored with a LUN. Attributes are available to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

The LUN attributes REST API allows you to create, update, delete, and discover attributes for a LUN. The LUN REST API also allows you to set attributes when a LUN is first created.

### Examples

## Retrieving all attributes from a LUN

This example uses the LUN attribute REST API endpoint with the `fields` query parameter to request the names and values.

```
# The API:
GET /api/storage/luns/{lun.uuid}/attributes

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes?fields=*' -H 'Accept: application/hal+json'

# The response:
{
"records": [
  {
    "name": "name1",
    "value": "value1",
    "_links": {
      "self": {
        "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name1"
      }
    }
  },
  {
    "name": "name2",
    "value": "value2",
    "_links": {
      "self": {
        "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes"
  }
}
}
```

This example uses the LUN REST endpoint with the `fields` query parameter to request the attributes properties.

```
# The API:
GET /api/storage/luns/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90?fields=attributes' -H 'Accept: application/hal+json'

# The response:
{
"uuid": "4bc204df-ecd8-4f35-8207-d0ccb4db3a90",
"name": "/vol/vol1/lun1",
"attributes": [
  {
    "name": "name1",
    "value": "name1",
    "_links": {
      "self": {
        "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name1"
      }
    }
  },
  {
    "name": "name2",
    "value": "value2",
    "_links": {
      "self": {
        "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name2"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-d0ccb4db3a90"
  }
}
}
```

### Adding an attribute to a LUN

```
# The API:
POST /api/storage/luns/{lun.uuid}/attributes

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes?return_records=true' -H 'Accept:
application/hal+json' -d '{ "name": "name1", "value": "value1" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "name": "name1",
    "value": "value1",
    "_links": {
      "self": {
        "href": "/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name1"
      }
    }
  }
]
}
```

#### Modifying an attribute value for a LUN

```
# The API
PATCH /api/storage/luns/{lun.uuid}/attributes/{name}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/4bc204df-ecd8-4f35-8207-
d0ccb4db3a90/attributes/name1' -H 'Accept: application/hal+json' -d '{
"value": "newValue" }'
```

#### Deleting an attribute from a LUN

```
# The API:
DELETE /api/storage/luns/{lun.uuid}/attributes/{name}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/storage/luns/4bc204df-ecd8-4f35-
8207-d0ccb4db3a90/attributes/name1' -H 'Accept: application/hal+json'
```

### Retrieve LUN attributes

GET `/storage/luns/{lun.uuid}/attributes`

**Introduced In:** 9.10

Retrieves LUN attributes.

#### Learn more

- DOC /storage/luns/{lun.uuid}/attributes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| value | string | query | False | Filter by value<br>• maxLength: 4091<br>• minLength: 1 |
| name | string | query | False | Filter by name<br>• maxLength: 4091<br>• minLength: 1 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | The number of records in the response. |
| records | array[lun_attribute] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "lun": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "name1",
      "value": "value1"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

lun

The LUN for which the attribute is set.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the LUN. |

lun_attribute

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| lun | lun | The LUN for which the attribute is set. |
| name | string | The attribute name. Required in POST. |
| value | string | The attribute value. Required in POST; valid in PATCH. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Add an attribute to a LUN

POST `/storage/luns/{lun.uuid}/attributes`

**Introduced In:** 9.10

Adds an attribute to a LUN.

#### Required properties

- `name` - The name of the attribute to add.
- `value` - The value of the attribute to add.

#### Learn more

- DOC /storage/luns/{lun.uuid}/attributes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| lun | lun | The LUN for which the attribute is set. |
| name | string | The attribute name. Required in POST. |

| Name | Type | Description |
| --- | --- | --- |
| value | string | The attribute value. Required in POST; valid in PATCH. |

### Example request

```
{
  "lun": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "name1",
  "value": "value1"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5374875 | The specified LUN was not found. |
| 5374928 | An incomplete attribute name/value pair was supplied. |
| 5374929 | The combined sizes of an attribute name and value are too large. |
| 5374930 | The attribute already exists for the LUN. |
| 5375060 | The copy start operation failed. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

lun

The LUN for which the attribute is set.

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the LUN. |

lun_attribute

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

| Name | Type | Description |
|---|---|---|
| lun | lun | The LUN for which the attribute is set. |
| name | string | The attribute name. Required in POST. |
| value | string | The attribute value. Required in POST; valid in PATCH. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a LUN attribute

DELETE `/storage/luns/{lun.uuid}/attributes/{name}`

**Introduced In:** 9.10

Deletes a LUN attribute.

#### Learn more

- DOC /storage/luns/{lun.uuid}/attributes

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| lun.uuid | string | path | True | |
| name | string | path | True | |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5374875 | The specified LUN was not found. |
| 5374931 | The specified attribute was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a LUN attribute

GET `/storage/luns/{lun.uuid}/attributes/{name}`

**Introduced In:** 9.10

Retrieves a LUN attribute.

#### Learn more

- DOC /storage/luns/{lun.uuid}/attributes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| name | string | path | True | The name of the attribute. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| lun | lun | The LUN for which the attribute is set. |
| name | string | The attribute name. Required in POST. |
| value | string | The attribute value. Required in POST; valid in PATCH. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "name1",
  "value": "value1"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5374875 | The specified LUN was not found. |
| 5374931 | The specified attribute was not found. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

lun

The LUN for which the attribute is set.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The unique identifier of the LUN. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a LUN attribute value

PATCH `/storage/luns/{lun.uuid}/attributes/{name}`

**Introduced In:** 9.10

Updates a LUN attribute value.

#### Learn more

- DOC /storage/luns/{lun.uuid}/attributes

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| name | string | path | True | The name of the attribute. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| lun | lun | The LUN for which the attribute is set. |
| value | string | The attribute value. Required in POST; valid in PATCH. |

### Example request

```
{
  "lun": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "value": "value1"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5374875 | The specified LUN was not found. |
| 5374929 | The combined sizes of an attribute name and value are too large. |
| 5374931 | The specified attribute was not found. |
| 5375060 | The move start operation failed. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

lun

The LUN for which the attribute is set.

| Name | Type | Description |
|---|---|---|
| uuid | string | The unique identifier of the LUN. |

lun_attribute

A name/value pair optionally stored with the LUN. Attributes are available to callers to persist small amounts of application-specific metadata. They are in no way interpreted by ONTAP.

Attribute names and values must be at least one byte and no more than 4091 bytes in length. The sum of the name and value lengths must be no more than 4092 bytes.

| Name | Type | Description |
|---|---|---|
| lun | lun | The LUN for which the attribute is set. |
| value | string | The attribute value. Required in POST; valid in PATCH. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve historical performance metrics for a LUN

GET `/storage/luns/{lun.uuid}/metrics`

**Introduced In:** 9.14

Retrieves historical space and performance metrics for a LUN.

### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| lun.uuid | string | path | True | The unique identifier of the LUN. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |
| status | string | query | False | Filter by status |
| timestamp | string | query | False | Filter by timestamp |
| throughput.other | integer | query | False | Filter by throughput.other |
| throughput.total | integer | query | False | Filter by throughput.total |
| throughput.write | integer | query | False | Filter by throughput.write |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| throughput.read | integer | query | False | Filter by throughput.read |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| uuid | string | query | False | Filter by uuid |
| duration | string | query | False | Filter by duration |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8585947 | No metrics are available for the requested object. |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |
| uuid | string | The unique identifier of the LUN. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve historical performance metrics for a LUN for a specific time

GET `/storage/luns/{lun.uuid}/metrics/{timestamp}`

**Introduced In:** 9.14

Retrieves historical space and performance metrics for a LUN for a specific time.

### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| lun.uuid | string | path | True | The unique identifier of the LUN. |
| timestamp | string | path | True | The timestamp of the space and performance data.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the LUN. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 8585947 | No metrics are available for the requested object. |

| Error Code | Description |
| --- | --- |
| 8586225 | An unexpected error occurred retrieving metrics for the requested object. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
|---|---|---|
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

# SVM

## ONTAP REST API SVM endpoints

### Overview

Storage Virtual Machine (SVM) APIs enable you to manage SVMs and their attributes, including the configuration of the CIFS and NFS protocols, export policies, name mappings between CIFS and NFS users, and network services.

SVMs contain data volumes and one or more network interfaces through which they serve data to the clients. SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain, and each SVM can be managed independently by its SVM administrator.

In a cluster, SVMs facilitate data access. A cluster must have at least one SVM to serve data. Multiple SVMs can coexist in a single cluster without being bound to any particular node in the cluster. However, they are bound to the physical cluster on which they exist.

SVMs with volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts. The volumes within each NAS SVM are related to each other through junctions and they are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

The ONTAP REST APIs only expose data SVMs as an SVM. The information and configuration associated with the cluster and nodes are exposed from REST, but the cluster and nodes are not treated as if they are a type of SVM from REST APIs. Some APIs that expose both cluster-owned resources and SVM-owned resources from the same endpoint only return and support the "svm" sub-object for the resources that are within a data SVM. In those endpoints, the resources that are not in a data SVM do not return the "svm" sub-object. Generally, such endpoints have a "scope" attribute that returns either "svm" or "cluster" to identify the resource as either a cluster-level resource or one that is completely contained in a data SVM.

## Manage SVM migration

### Manage SVM migration

### Overview

You can use this API to migrate an SVM from a source cluster to a destination cluster. During the migration, you can pause, resume, or abort the operation, and retrieve the migration status and transfer status of volumes in the SVM.
The migrations APIs are only accessible from the destination cluster.

#### Precondition

The cluster peering relationship should be created between the source and destination clusters prior to using this API.

#### SVM migration APIs

The following APIs are used to manage SVM migration:

- POST /api/svm/migrations
- PATCH /api/svm/migrations/{uuid}
- GET /api/svm/migrations/
- GET /api/svm/migrations/{uuid}
- GET /api/svm/migrations/{svm_migration.uuid}/volumes
- GET /api/svm/migrations/{svm_migration.uuid}/volumes/{volume.uuid}
- DELETE /api/svm/migrations/{uuid}

#### Important notes

The migration of temperature-sensitive storage efficiency (TSSE) volumes from AFF to FAS systems has a known limitation. The migration preserves the TSSE savings and after migration completes, the volume will not receive any further TSSE savings for the new data written on the volume.

#### Starting a migration

To start a migration operation, issue a POST request to /svm/migrations. Parameters are provided in the body of the POST request to specify the source cluster and the source SVM.

#### Monitoring the status of the SVM migration

You can use GET /svm/migrations to retrieve the status of the SVM migration and GET /svm/migrations/{svm_migration.uuid}/volumes to retrieve the transfer status of the volumes in the SVM migration.

### Possible errors before starting the migration

Configurations in the POST /svm/migrations request are validated before the SVM migration starts. If an invalid configuration is found or the migration pre-checks fail, an HTTP error code in the 4xx range is returned. No SVM migration operation is started.

### Polling the migration operation

After a successful POST /svm/migrations request is issued, an HTTP error code of 202 is returned along with a migration UUID and link in the body of the response. The SVM migration continues asynchronously and is monitored using the migration UUID and the GET /svm/migrations/{uuid} API.

### Errors during the migration operation

If a failure occurs during the SVM migration, the GET /svm/migrations response provides details of the error along with any error code fields.

### Pausing the migration operation

You can use PATCH /svm/migrations/{uuid} with the action "pause" to pause the SVM migration to update the SVM configuration on the source SVM.

### Resuming the migration operation

You can use PATCH /svm/migrations{uuid} with the action "resume" to resume the SVM migration from a paused state.
You can modify the throttle value when you resume the SVM migration. To set the throttle value to unlimited, specify the throttle value as 0.

### Aborting the migration operation

You can use DELETE /svm/migrations/{uuid} to delete the SVM on the destination cluster if the SVM migration has failed or is paused.
Use the DELETE /svm/migrations/{uuid} request to remove the SVM on the source cluster; this might be used when communication between the source and destination cluster is reduced.

#### Retrieving the migration status

You can use GET /svm/migrations/{uuid} to retrieve the current status of your migration.

#### Retrieving the volume transfer status

You can use GET /svm/migrations/{svm_migration.uuid}/volumes to retrieve the current transfer status of all volumes in the migrating SVM.

#### Retrieving the volume transfer status of a specific volume

You can use GET /svm/migrations/{svm_migrations.uuid}/volumes/{volume.uuid} to retrieve the transfer status of a specific volume in the migrating SVM.

---

#### Examples

### Starting a Migration

```
# API
/api/svm/migrations
```

### POST body included from file

```
svm_migration_post_body.txt:
{
"source": {
  "svm": {
    "name": "vs1"
  },
  "cluster": {
    "name": "siteB"
  }
}
}
curl -X POST https://<mgmt-ip>/api/svm/migrations -H "Content-Type:
application+hal/json" -d "@svm_migration_post_body.txt"
```

**Inline POST body**

```
curl -X POST https://<mgmt-ip>/api/svm/migrations -H "Content-Type:
application+hal/json" -d '{"source": {"svm": {"name": "vs1"}, "cluster":
{"name": "siteB"}}}'
```

**POST Response**

```
Date: Wed, 25 Aug 2021 19:04:47 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/svm/migrations/517c5e74-05d7-11ec-a40f-005056bba9a5
Content-Length: 189
Content-Type: application/hal+json
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"job": {
  "uuid": "5184a3e1-05d7-11ec-a40f-005056bba9a5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/5184a3e1-05d7-11ec-a40f-005056bba9a5"
    }
  }
}
}
```

### Retrieving POST Job status

Use the link provided in the response to the POST request to fetch the status of the start operation

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/5184a3e1-05d7-11ec-a40f-
005056bba9a5
```

#### Response

```
Date: Wed, 25 Aug 2021 19:05:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 224
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"uuid": "5184a3e1-05d7-11ec-a40f-005056bba9a5",
"description": "POST /api/svm/migrations/517c5e74-05d7-11ec-a40f-
005056bba9a5",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-08-25T15:04:48-04:00",
"end_time": "2021-08-25T15:04:57-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/5184a3e1-05d7-11ec-a40f-005056bba9a5"
  }
}
}
```

### Retrieving all migrations

The location header in the POST operation provides the uuid of the migrate operation that was started using POST. Also, you can list all the migrate operations using the collection GET.

#### Request

```
curl -X GET https://<mgmt-ip>/api/svm/migrations
```

**Response**

```
Date: Wed, 25 Aug 2021 19:05:11 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 170
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"records": [
  {
    "uuid": "517c5e74-05d7-11ec-a40f-005056bba9a5",
    "_links": {
      "self": {
        "href": "/api/svm/migrations/517c5e74-05d7-11ec-a40f-005056bba9a5"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/svm/migrations/"
  }
}
}
```

**Retrieving a specific migration**

**Request**

```
curl -X GET https://<mgmt-ip>/api/svm/migrations/517c5e74-05d7-11ec-a40f-
005056bba9a5
```

**Response**

```
Date: Wed, 25 Aug 2021 19:05:33 GMT
```

```
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 379
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"uuid": "517c5e74-05d7-11ec-a40f-005056bba9a5",
"state": "setup_configuration",
"start_time": "2021-08-25T15:04:49-04:00",
"current_operation": "start",
"source": {
  "svm": {
    "uuid": "424b6002-fb1a-11eb-9383-005056bbcf32",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/424b6002-fb1a-11eb-9383-005056bbcf32"
      }
    }
  },
  "cluster": {
    "uuid": "b54babec-fb14-11eb-9383-005056bbcf32",
    "name": "siteB",
    "_links": {
      "self": {
        "href": "/api/cluster/peers/b54babec-fb14-11eb-9383-005056bbcf32"
      }
    }
  }
},
"destination": {
  "ipspace": {
    "uuid": "f305cf0b-fb14-11eb-829d-005056bba9a5",
    "name": "Default"
  }
},
"auto_cutover": false,
"auto_source_cleanup": false,
"throttle": 0,
"_links": {
  "self": {
    "href": "/api/svm/migrations/517c5e74-05d7-11ec-a40f-005056bba9a5"
```

```
  }
}
}
```

#### Pausing a migration

To pause the migration use the PATCH request on the migration UUID.

#### Request

```
curl -X PATCH https://<mgmt-ip>/api/svm/migrations/517c5e74-05d7-11ec-
a40f-005056bba9a5?action=pause
```

#### Response

```
Date: Wed, 25 Aug 2021 19:06:11 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"job": {
  "uuid": "82dea7c7-05d7-11ec-a40f-005056bba9a5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/82dea7c7-05d7-11ec-a40f-005056bba9a5"
    }
  }
}
}
```

#### Monitoring PATCH job status

Use the link provided in the response of the PATCH request to fetch the information of the patch job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/82dea7c7-05d7-11ec-a40f-
005056bba9a5
```

### Response

```
Date: Wed, 25 Aug 2021 21:40:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 222
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"uuid": "82dea7c7-05d7-11ec-a40f-005056bba9a5",
"description": "PATCH /api/svm/migrations/517c5e74-05d7-11ec-a40f-
005056bba9a5",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-08-25T15:06:11-04:00",
"end_time": "2021-08-25T15:06:11-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/82dea7c7-05d7-11ec-a40f-005056bba9a5"
  }
}
}
```

### Aborting a migration

To abort the migration use the DELETE request on the migration UUID.

### Request

```
curl -X DELETE https://<mgmt-ip>/api/svm/migrations/517c5e74-05d7-11ec-
a40f-005056bba9a5
```

### Response

```
Date: Wed, 25 Aug 2021 22:57:23 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"job": {
  "uuid": "cf870f19-05f7-11ec-a40f-005056bba9a5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/cf870f19-05f7-11ec-a40f-005056bba9a5"
    }
  }
}
}
```

#### Monitoring DELETE job status

Use the link provided in the response of the PATCH request to fetch the information of the patch job.

#### Request

```
curl -X GET https://<mgmt-ip>/api/cluster/jobs/cf870f19-05f7-11ec-a40f-
005056bba9a5
```

#### Response

```
Date: Wed, 25 Aug 2021 23:05:47 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 228
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
{
"uuid": "cf870f19-05f7-11ec-a40f-005056bba9a5",
"description": "DELETE /api/svm/migrations/517c5e74-05d7-11ec-a40f-
005056bba9a5",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-08-25T18:57:23-04:00",
"end_time": "2021-08-25T18:57:24-04:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/cf870f19-05f7-11ec-a40f-005056bba9a5"
  }
}
}
```

### Retrieve the SVM migration status

GET `/svm/migrations`

**Introduced In:** 9.10

Retrieves the SVM migration status.

#### Related ONTAP commands

- `vserver migrate show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| post_ponr_retry_count | integer | query | False | Filter by post_ponr_retry_count<br>• Introduced in: 9.17 |
| auto_source_cleanup | boolean | query | False | Filter by auto_source_cleanup |
| source.cluster.name | string | query | False | Filter by source.cluster.name |
| source.cluster.uuid | string | query | False | Filter by source.cluster.uuid |
| source.svm.name | string | query | False | Filter by source.svm.name |
| source.svm.uuid | string | query | False | Filter by source.svm.uuid |
| point_of_no_return | boolean | query | False | Filter by point_of_no_return |
| destination.ipspace.name | string | query | False | Filter by destination.ipspace.name |
| destination.ipspace.uuid | string | query | False | Filter by destination.ipspace.uuid |
| uuid | string | query | False | Filter by uuid |
| current_operation | string | query | False | Filter by current_operation |
| messages.message | string | query | False | Filter by messages.message |
| messages.code | string | query | False | Filter by messages.code |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| auto_cutover | boolean | query | False | Filter by auto_cutover |
| restart_count | integer | query | False | Filter by restart_count |
| throttle | integer | query | False | Filter by throttle<br>• Introduced in: 9.12 |
| last_operation | string | query | False | Filter by last_operation |
| time_metrics.start_time | string | query | False | Filter by time_metrics.start_time |
| time_metrics.last_resume_time | string | query | False | Filter by time_metrics.last_resume_time |
| time_metrics.last_pause_time | string | query | False | Filter by time_metrics.last_pause_time |
| time_metrics.last_post_ponr_retry_time | string | query | False | Filter by time_metrics.last_post_ponr_retry_time<br>• Introduced in: 9.17 |
| time_metrics.cutover_start_time | string | query | False | Filter by time_metrics.cutover_start_time |
| time_metrics.cutover_trigger_time | string | query | False | Filter by time_metrics.cutover_trigger_time |
| time_metrics.cutover_complete_time | string | query | False | Filter by time_metrics.cutover_complete_time |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| time_metrics.end_time | string | query | False | Filter by time_metrics.end_time |
| last_failed_state | string | query | False | Filter by last_failed_state |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[svm_migration] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "check_only": true,
      "current_operation": "string",
      "destination": {
        "ipspace": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "Default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "last_failed_state": "string",
      "last_operation": "string",
      "messages": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "post_ponr_retry_count": 0,
      "restart_count": 0,
      "source": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "state": "string",
      "throttle": 0,
      "time_metrics": {
        "cutover_complete_time": "2020-12-02 19:30:19 -0800",
        "cutover_start_time": "2020-12-02 18:20:19 -0800",
        "cutover_trigger_time": "2020-12-02 19:15:19 -0800",
        "end_time": "2020-12-02 19:36:19 -0800",
        "last_pause_time": "2020-12-02 18:50:19 -0800",
        "last_post_ponr_retry_time": "2020-12-02 19:30:19 -0800",
        "last_resume_time": "2020-12-02 18:54:19 -0800",
        "start_time": "2020-12-02 18:36:19 -0800"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

aggregate

Aggregate to use for volume creation.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

volume

Property indicating the source volume.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

volume_aggregate_pairs

Volume-aggregate pair information.

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate to use for volume creation. |
| volume | volume | Property indicating the source volume. |

volume_placement

Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Optional property used to specify the list of desired aggregates to use for volume creation in the destination. |
| volume_aggregate_pairs | array[volume_aggregate_pairs] | Optional property used to specify the list of desired volume-aggregate pairs in the destination. |

destination

Destination cluster details for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

interface

Network interface on the source SVM.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

port

Port to use for IP interface placement on the destination SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| node | node | |
| uuid | string | |

ip_interfaces

IP interface and network port pair information.

| Name | Type | Description |
|---|---|---|
| interface | interface | Network interface on the source SVM. |
| port | port | Port to use for IP interface placement on the destination SVM. |

svm_migration_ip_interface_placement

Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided.

| Name | Type | Description |
|---|---|---|
| ip_interfaces | array[ip_interfaces] | List of source SVM's IP interface and port pairs on the destination for migrating the source SVM's IP interfaces. |

messages

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

cluster

Source cluster for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

Source SVM

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

source

Source cluster details for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster | Source cluster for the SVM migration. |
| svm | svm | Source SVM |

time_metrics

Various time metrics details

| Name | Type | Description |
| --- | --- | --- |
| cutover_complete_time | string | Cutover end time |

| Name | Type | Description |
| --- | --- | --- |
| cutover_start_time | string | Cutover start time |
| cutover_trigger_time | string | Cutover trigger time |
| end_time | string | Migration end time |
| last_pause_time | string | Last migration pause time |
| last_post_ponr_retry_time | string | Last post point of no return retry time |
| last_resume_time | string | Last migration resume time |
| start_time | string | Migration start time |

svm_migration

Provides information on SVM migration, default and user specified configurations, the state of the migration, and volume transfer metrics.

| Name | Type | Description |
| --- | --- | --- |
| auto_cutover | boolean | Optional property that when set to true automatically performs cutover when the migration state reaches "ready for cutover". |
| auto_source_cleanup | boolean | Optional property that when set to true automatically cleans up the SVM on the source cluster after the migration cutover. |
| current_operation | string | |
| destination | destination | Destination cluster details for the SVM migration. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |

| Name | Type | Description |
| --- | --- | --- |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an SVM migration operation

POST `/svm/migrations`

**Introduced In:** 9.10

Creates an SVM migration operation. This API must be executed on the destination cluster. This API creates an SVM on the destination cluster and preserves the SVM's identity specified in the source cluster. Optionally, you can specify the aggregate list for creating the volumes, and IPspace. You can perform pre-checks to verify if SVM migration is possible, by setting the "check-only" option to "true". By default the values for auto-source-cleanup and auto-cutover is true.

#### Required properties

- `source.svm.name` or `source.svm.uuid` - Source SVM name or source SVM UUID.
- `source.cluster.name` or `source.cluster.uuid` - Source cluster name or source cluster UUID

#### Optional properties

- `destination.ipspace.name` or `destination.ipspace.uuid` - Destination IP Space name or UUID where the SVM will be migrated to.
- `destination.volume_placement.aggregates` - List of aggregates where the migrating volumes should go on the destination.
- `destination.volume_placement.volume_aggregate_pairs` - List of volume aggregate pairs indicating where the migrating volumes should go on the destination.
- `ip_interface_placement` - List of source SVM's IP interface and port pairs on the destination for migrating the SVM's IP interfaces.
- `auto_cutover` - Option to specify whether to perform cutover automatically. Default is true.
- `auto_source_cleanup` - Option to specify whether to perform source cleanup automatically. Default is true.
- `check_only` - Option to perform all the prechecks for migrate without actually starting the migrate. Default is false.
- `throttle` - Option to specify a throttle value in KB/s. Defaults to unlimited.

#### Related ONTAP commands

- `vserver migrate start`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| auto_cutover | boolean | Optional property that when set to true automatically performs cutover when the migration state reaches "ready for cutover". |
| auto_source_cleanup | boolean | Optional property that when set to true automatically cleans up the SVM on the source cluster after the migration cutover. |
| check_only | boolean | Optional property that when set to true performs only migration pre-checks not the actual migration. |
| current_operation | string | |
| destination | destination | Destination cluster details for the SVM migration. |
| ip_interface_placement | svm_migration_ip_interface_placement | Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |

| Name | Type | Description |
| --- | --- | --- |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

**Example request**


```
{
  "current_operation": "string",
  "destination": {
    "ipspace": {
      "name": "Default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "volume_placement": {
      "aggregates": [
        {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "volume_aggregate_pairs": [
        {
          "aggregate": {
            "name": "aggr1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "volume": {
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ]
    }
  },
  "ip_interface_placement": {
    "ip_interfaces": [
      {
        "interface": {
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "port": {
          "name": "e1b",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
```

```
  },
  "last_failed_state": "string",
  "last_operation": "string",
  "messages": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "post_ponr_retry_count": 0,
  "restart_count": 0,
  "source": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "string",
  "throttle": 0,
  "time_metrics": {
    "cutover_complete_time": "2020-12-02 19:30:19 -0800",
    "cutover_start_time": "2020-12-02 18:20:19 -0800",
    "cutover_trigger_time": "2020-12-02 19:15:19 -0800",
    "end_time": "2020-12-02 19:36:19 -0800",
    "last_pause_time": "2020-12-02 18:50:19 -0800",
    "last_post_ponr_retry_time": "2020-12-02 19:30:19 -0800",
    "last_resume_time": "2020-12-02 18:54:19 -0800",
    "start_time": "2020-12-02 18:36:19 -0800"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[svm_migration] | |

**Example response**


```
{
  "job": {
    "uuid": "string"
  },
  "num_records": 1,
  "records": [
    {
      "current_operation": "string",
      "destination": {
        "ipspace": {
          "name": "Default",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "volume_placement": {
          "aggregates": [
            {
              "name": "aggr1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          ],
          "volume_aggregate_pairs": [
            {
              "aggregate": {
                "name": "aggr1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              },
              "volume": {
                "name": "volume1",
                "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
              }
            }
          ]
        }
      },
      "ip_interface_placement": {
        "ip_interfaces": [
          {
            "interface": {
              "name": "lif1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "port": {
              "name": "e1b",
              "node": {
```

```
                "name": "node1"
              },
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      },
      "last_failed_state": "string",
      "last_operation": "string",
      "messages": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "post_ponr_retry_count": 0,
      "restart_count": 0,
      "source": {
        "cluster": {
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "state": "string",
      "throttle": 0,
      "time_metrics": {
        "cutover_complete_time": "2020-12-02 19:30:19 -0800",
        "cutover_start_time": "2020-12-02 18:20:19 -0800",
        "cutover_trigger_time": "2020-12-02 19:15:19 -0800",
        "end_time": "2020-12-02 19:36:19 -0800",
        "last_pause_time": "2020-12-02 18:50:19 -0800",
        "last_post_ponr_retry_time": "2020-12-02 19:30:19 -0800",
        "last_resume_time": "2020-12-02 18:54:19 -0800",
        "start_time": "2020-12-02 18:36:19 -0800"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262245 | The value provided is invalid. |
| 13172746 | SVM migration cannot be started. This is a generic code, see the response message for details. |
| 13173748 | Migrate request cannot contain both \"aggregates\" and \"volume_aggregate_pairs\" within the \"volume_placement\" object. |
| 13173758 | The property \"{property}\" is not supported for this operation. |
| 13173759 | Migrate operation failed. To use LIF placement specify either the port UUID, or both the port name and the port node name. |
| 13173760 | Migrate operation failed. LIF placement requires either the port node name or the port UUID to be specified if the port name is specified. |
| 13173761 | Migrate operation failed. LIF placement requires either the port name or the port UUID to be specified if the port node name is specified. |
| 13173762 | Migrate operation failed. To use LIF placement specify at least one of the following: IP interface UUID, IP interface name, or IP interface IP. |
| 13173764 | Migrate operation failed because LIF placement is not supported on the destination cluster. Both clusters must have an effective cluster version of 9.12.1 or later. |
| 13173765 | Migrate operation failed. Unable to determine if LIF placement is supported. Reason: \"{Reason}\" |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

aggregate

Aggregate to use for volume creation.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

volume

Property indicating the source volume.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

volume_aggregate_pairs

Volume-aggregate pair information.

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate to use for volume creation. |
| volume | volume | Property indicating the source volume. |

volume_placement

Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Optional property used to specify the list of desired aggregates to use for volume creation in the destination. |
| volume_aggregate_pairs | array[volume_aggregate_pairs] | Optional property used to specify the list of desired volume-aggregate pairs in the destination. |

destination

Destination cluster details for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| ipspace | ipspace | Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used. |
| volume_placement | volume_placement | Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state. |

ip

IP information

interface

Network interface on the source SVM.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port to use for IP interface placement on the destination SVM.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| node | node | |
| uuid | string | |

ip_interfaces

IP interface and network port pair information.

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | Network interface on the source SVM. |
| port | port | Port to use for IP interface placement on the destination SVM. |

svm_migration_ip_interface_placement

Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided.

| Name | Type | Description |
| --- | --- | --- |
| ip_interfaces | array[ip_interfaces] | List of source SVM's IP interface and port pairs on the destination for migrating the source SVM's IP interfaces. |

messages

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

cluster

Source cluster for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

svm

Source SVM

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

source

Source cluster details for the SVM migration.

| Name | Type | Description |
|---|---|---|
| cluster | cluster | Source cluster for the SVM migration. |
| svm | svm | Source SVM |

time_metrics

Various time metrics details

| Name | Type | Description |
|---|---|---|
| cutover_complete_time | string | Cutover end time |
| cutover_start_time | string | Cutover start time |
| cutover_trigger_time | string | Cutover trigger time |
| end_time | string | Migration end time |
| last_pause_time | string | Last migration pause time |
| last_post_ponr_retry_time | string | Last post point of no return retry time |
| last_resume_time | string | Last migration resume time |
| start_time | string | Migration start time |

svm_migration

Provides information on SVM migration, default and user specified configurations, the state of the

migration, and volume transfer metrics.

| Name | Type | Description |
|---|---|---|
| auto_cutover | boolean | Optional property that when set to true automatically performs cutover when the migration state reaches "ready for cutover". |
| auto_source_cleanup | boolean | Optional property that when set to true automatically cleans up the SVM on the source cluster after the migration cutover. |
| check_only | boolean | Optional property that when set to true performs only migration pre-checks not the actual migration. |
| current_operation | string | |
| destination | destination | Destination cluster details for the SVM migration. |
| ip_interface_placement | svm_migration_ip_interface_placement | Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |

| Name | Type | Description |
| --- | --- | --- |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the transfer status of volumes in an SVM

GET `/svm/migrations/{svm_migration.uuid}/volumes`

**Introduced In:** 9.10

Retrieves the transfer status of the volumes in the SVM.

#### Related ONTAP commands

- `vserver migrate show-volume`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm_migration.uuid | string | path | True | Migration UUID |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[svm_migration_volume] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "errors": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "transfer_state": "string",
      "volume": {
        "_links": {
```

```
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

errors

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Node in the destination cluster where the volume is hosted

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM information

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume information in the destination cluster

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

svm_migration_volume

Volume transfer information

| Name | Type | Description |
| --- | --- | --- |
| _links | self_link | |
| errors | array[errors] | List of transfer errors |
| healthy | boolean | Indicates whether the volume transfer relationship is healthy. |
| node | node | Node in the destination cluster where the volume is hosted |

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM information |
| transfer_state | string | Status of the transfer. |
| volume | volume | Volume information in the destination cluster |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the volume transfer status for a volume

GET `/svm/migrations/{svm_migration.uuid}/volumes/{volume.uuid}`

**Introduced In:** 9.10

Retrieves the volume transfer status of the specified volume.uuid.

#### Related ONTAP commands

- `vserver migrate show-volume`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm_migration.uuid | string | path | True | Migration UUID |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| volume.uuid | string | path | True | Volume UUID |
| errors.message | string | query | False | Filter by errors.message |
| errors.code | string | query | False | Filter by errors.code |
| volume.name | string | query | False | Filter by volume.name |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| healthy | boolean | query | False | Filter by healthy |
| transfer_state | string | query | False | Filter by transfer_state |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| errors | array[errors] | List of transfer errors |
| healthy | boolean | Indicates whether the volume transfer relationship is healthy. |
| node | node | Node in the destination cluster where the volume is hosted |
| svm | svm | SVM information |
| transfer_state | string | Status of the transfer. |
| volume | volume | Volume information in the destination cluster |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "errors": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "transfer_state": "string",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

errors

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Node in the destination cluster where the volume is hosted

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM information

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

Volume information in the destination cluster

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Delete SVM migrations

DELETE `/svm/migrations/{uuid}`

**Introduced In:** 9.10

Deletes the SVM migration.

#### Related ONTAP commands

- `vserver migrate abort`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | SVM migration UUID. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |
| 13173738 | REST API DELETE method \"/api/svm/migrations\" is only supported on the destination cluster. Issue the REST API DELETE request to the destination cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve the migration status for an SVM**

GET `/svm/migrations/{uuid}`

**Introduced In:** 9.10

Retrieves the migration status of an individual SVM.

#### Important notes

- The "migrations" object includes a large set of fields and can be expensive to retrieve.
- REST APIs only expose a data SVM as an SVM.
- There are subsystem specific errors that can be returned from this endpoint. If a subsystem specific error is returned and this is the first migrate operation attempt, it is embedded in one of the following errors. If a subsystem specific error is returned and this is not the first migrate operation attempt, the subsystem specific error is returned directly.

#### Example

Retrieving an individual SVM migration status.

```
  GET "/api/svm/migrations/a14ae39f-8d85-11e9-b4a7-
00505682dc8b/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Migration UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| auto_cutover | boolean | Optional property that when set to true automatically performs cutover when the migration state reaches "ready for cutover". |
| auto_source_cleanup | boolean | Optional property that when set to true automatically cleans up the SVM on the source cluster after the migration cutover. |
| current_operation | string | |

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | Destination cluster details for the SVM migration. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

**Example response**


```
{
  "check_only": true,
  "current_operation": "string",
  "destination": {
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "Default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "last_failed_state": "string",
  "last_operation": "string",
  "messages": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "post_ponr_retry_count": 0,
  "restart_count": 0,
  "source": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
```

```
  "state": "string",
  "throttle": 0,
  "time_metrics": {
    "cutover_complete_time": "2020-12-02 19:30:19 -0800",
    "cutover_start_time": "2020-12-02 18:20:19 -0800",
    "cutover_trigger_time": "2020-12-02 19:15:19 -0800",
    "end_time": "2020-12-02 19:36:19 -0800",
    "last_pause_time": "2020-12-02 18:50:19 -0800",
    "last_post_ponr_retry_time": "2020-12-02 19:30:19 -0800",
    "last_resume_time": "2020-12-02 18:54:19 -0800",
    "start_time": "2020-12-02 18:36:19 -0800"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |
| 13173739 | Migrate pause operation failed. Retry pause operation using REST API PATCH method \"/api/svm/migrations/<migration_uuid>?action=pause\\". Reason: {Reason}</migration_uuid> |
| 13173740 | Migrate abort operation failed. Retry abort operation by using REST API DELETE method \"/api/svm/migrations/<migration_uuid>\\". Reason: {Reason}</migration_uuid> |
| 13173741 | Migrate failed. Retry the migrate by running the resume operation using REST API PATCH method \"/api/svm/migrations/<migration_uuid>?action=resume\\". Reason: {Reason}</migration_uuid> |
| 13173742 | Migrate operation status: {Reason}. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

aggregate

Aggregate to use for volume creation.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

volume

Property indicating the source volume.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

volume_aggregate_pairs

Volume-aggregate pair information.

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate to use for volume creation. |
| volume | volume | Property indicating the source volume. |

volume_placement

Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state.

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Optional property used to specify the list of desired aggregates to use for volume creation in the destination. |
| volume_aggregate_pairs | array[volume_aggregate_pairs] | Optional property used to specify the list of desired volume-aggregate pairs in the destination. |

destination

Destination cluster details for the SVM migration.

| Name | Type | Description |
|---|---|---|
| ipspace | ipspace | Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used. |

ip

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |

interface

Network interface on the source SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
|---|---|---|
| name | string | Name of node on which the port is located. |

port

Port to use for IP interface placement on the destination SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

ip_interfaces

IP interface and network port pair information.

| Name | Type | Description |
| --- | --- | --- |
| interface | interface | Network interface on the source SVM. |
| port | port | Port to use for IP interface placement on the destination SVM. |

svm_migration_ip_interface_placement

Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided.

| Name | Type | Description |
| --- | --- | --- |
| ip_interfaces | array[ip_interfaces] | List of source SVM's IP interface and port pairs on the destination for migrating the source SVM's IP interfaces. |

messages

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

cluster

Source cluster for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

svm

Source SVM

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

source

Source cluster details for the SVM migration.

| Name | Type | Description |
|---|---|---|
| cluster | cluster | Source cluster for the SVM migration. |
| svm | svm | Source SVM |

time_metrics

Various time metrics details

| Name | Type | Description |
|---|---|---|
| cutover_complete_time | string | Cutover end time |
| cutover_start_time | string | Cutover start time |
| cutover_trigger_time | string | Cutover trigger time |
| end_time | string | Migration end time |
| last_pause_time | string | Last migration pause time |
| last_post_ponr_retry_time | string | Last post point of no return retry time |
| last_resume_time | string | Last migration resume time |
| start_time | string | Migration start time |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update actions performed during an SVM migration

PATCH `/svm/migrations/{uuid}`

**Introduced In:** 9.10

Actions that can be performed during an SVM migration.

#### Related ONTAP commands

- `vserver migrate pause`
- `vserver migrate resume`
- `vserver migrate cutover`
- `vserver migrate source-cleanup`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | SVM migration UUID |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| action | string | query | False | The pause action pauses the SVM migration. This action stops data transfer and configuration replication. This operation must be performed on the destination cluster. The resume action resumes an SVM migration from a paused or failed state. If the SVM migration is in a cleanup_failed state, volume placement is ignored. This operation must be performed on the destination cluster. The cutover action triggers the cutover of an SVM from the source cluster to the destination cluster. The source_clean up action performs a clean up of the SVM on the source cluster.<br><br>• enum: ["pause", "resume", "cutover", "source_cleanup "] |
| auto_cutover | boolean | query | False | Optional property that when set to true automatically performs cutover when the migration state reaches "ready for cutover".<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| auto_source_cleanup | boolean | query | False | Optional property that when set to true automatically cleans up the SVM on the source cluster after the migration cutover.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| current_operation | string | |

| Name | Type | Description |
| --- | --- | --- |
| destination | destination | Destination cluster details for the SVM migration. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

### Example request

```
{
  "auto_cutover": true,
  "auto_source_cleanup": true,
  "check_only": true,
  "current_operation": "string",
  "destination": {
    "volume_placement": {
      "volume_aggregate_pairs": [
        {
          "aggregate": {
            "name": "aggr1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "volume": {
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ]
    }
  },
  "last_failed_state": "string",
  "last_operation": "string",
  "messages": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "post_ponr_retry_count": 0,
  "restart_count": 0,
  "state": "string",
  "throttle": 0,
  "time_metrics": {
    "cutover_complete_time": "2020-12-02 19:30:19 -0800",
    "cutover_start_time": "2020-12-02 18:20:19 -0800",
    "cutover_trigger_time": "2020-12-02 19:15:19 -0800",
    "end_time": "2020-12-02 19:36:19 -0800",
    "last_pause_time": "2020-12-02 18:50:19 -0800",
    "last_post_ponr_retry_time": "2020-12-02 19:30:19 -0800",
    "last_resume_time": "2020-12-02 18:54:19 -0800",
    "start_time": "2020-12-02 18:36:19 -0800"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 13172783 | Migrate RDB lookup failed |
| 13173737 | REST API PATCH method \"/api/svm/migrations\" is only supported on the destination cluster. Issue the REST API PATCH request to the destination cluster. |
| 13173746 | Migrate resume operation failed. Cannot specify volume granular placement during resume if aggregate placement was specified during start operation. |
| 13173747 | Migrate operation failed. Volume placement can only be specified on PATCH with an action of \"resume\". |
| 13173748 | Migrate request cannot contain both \"aggregates\" and \"volume_aggregate_pairs\" within the \"volume_placement\" object. |
| 13173763 | Migrate operation failed. LIF placement is not supported in PATCH operations. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

Optional property used to specify which IPspace to use for the SVM. By default, the "default" ipspace is used.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

aggregate

Aggregate to use for volume creation.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

volume

Property indicating the source volume.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

volume_aggregate_pairs

Volume-aggregate pair information.

| Name | Type | Description |
| --- | --- | --- |
| aggregate | aggregate | Aggregate to use for volume creation. |
| volume | volume | Property indicating the source volume. |

volume_placement

Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state.

| Name | Type | Description |
| --- | --- | --- |
| volume_aggregate_pairs | array[volume_aggregate_pairs] | Optional property used to specify the list of desired volume-aggregate pairs in the destination. |

destination

Destination cluster details for the SVM migration.

| Name | Type | Description |
| --- | --- | --- |
| volume_placement | volume_placement | Optional property to specify the source volume placement in the destination. It is input only and won't be returned by a subsequent GET. Volume placement is ignored if the migration resumes from the cleanup_failed state. |

ip

IP information

interface

Network interface on the source SVM.

| Name | Type | Description |
| --- | --- | --- |
| ip | ip | IP information |
| name | string | The name of the interface. If only the name is provided, the SVM scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the interface. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of node on which the port is located. |

port

Port to use for IP interface placement on the destination SVM.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| node | node | |
| uuid | string | |

ip_interfaces

IP interface and network port pair information.

| Name | Type | Description |
|---|---|---|
| interface | interface | Network interface on the source SVM. |
| port | port | Port to use for IP interface placement on the destination SVM. |

svm_migration_ip_interface_placement

Optional property used to specify the list of source SVM's IP interface and network port pairs in the destination for migrating the source SVM IP interfaces. Note that the SVM migration does not perform any reachability checks on the IP interfaces provided.

| Name | Type | Description |
|---|---|---|
| ip_interfaces | array[ip_interfaces] | List of source SVM's IP interface and port pairs on the destination for migrating the source SVM's IP interfaces. |

messages

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

cluster

Source cluster for the SVM migration.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

Source SVM

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

source

Source cluster details for the SVM migration.

time_metrics

Various time metrics details

| Name | Type | Description |
| --- | --- | --- |
| cutover_complete_time | string | Cutover end time |
| cutover_start_time | string | Cutover start time |
| cutover_trigger_time | string | Cutover trigger time |
| end_time | string | Migration end time |
| last_pause_time | string | Last migration pause time |
| last_post_ponr_retry_time | string | Last post point of no return retry time |
| last_resume_time | string | Last migration resume time |
| start_time | string | Migration start time |

svm_migration

Provides information on SVM migration, default and user specified configurations, the state of the migration, and volume transfer metrics.

| Name | Type | Description |
| --- | --- | --- |
| current_operation | string | |
| destination | destination | Destination cluster details for the SVM migration. |
| last_failed_state | string | Indicates the state of the migration. |
| last_operation | string | |

| Name | Type | Description |
| --- | --- | --- |
| messages | array[messages] | Errors and warnings returned/displayed during migration. |
| point_of_no_return | boolean | Indicates if the migration has progressed beyond the point of no return. When true, the migration cannot be aborted or paused. When false, the migration can be paused or aborted. |
| post_ponr_retry_count | integer | Number of times the migration restarted after the point of no return. |
| restart_count | integer | Number of times migrate restarted the transfer, for example, rollback to transfer after starting the cutover. |
| source | source | Source cluster details for the SVM migration. |
| state | string | Indicates the state of the migration. |
| throttle | integer | Optional property to specify a throttle value in KB/s for each individual volume transfer. Defaults to 0 if not set, which is interpreted as unlimited. The minimum throttle value is 4 KB/s, so if you specify a throttle value between 1 and 4, it will be treated as if you specified 4. |
| time_metrics | time_metrics | Various time metrics details |
| uuid | string | SVM migration UUID |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SVM peer permissions

### Manage SVM peer permissions

#### Managing SVM peer permissions

A cluster administrator can provide permissions for use during intercluster SVM peer relationship creation. Once this permission exists for a local SVM and peer cluster combination on a local cluster, no explicit SVM peer accept (or REST PATCH) API is required for any incoming SVM peer relationship creation requests from a remote cluster for that local SVM. Peer relationship directly changes the state to peered on both clusters. Use an SVM name as "*" to create permissions that apply to all local SVMs.

### SVM peer permission APIs

The following APIs are used to manage SVM peer permissions:

- GET /api/svm/peer-permissions
- POST /api/svm/peer-permissions
- GET /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}
- PATCH /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}

- DELETE /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}

**Retrieve SVM peer permissions**

GET `/svm/peer-permissions`

**Introduced In:** 9.6

Retrieves the list of SVM peer permissions.

#### Related ONTAP commands

- `vserver peer permission show`

#### Examples

The following examples show how to retrieve a collection of SVM peer permissions based on a query.

1. Retrieves a list of SVM peer permissions of a specific local SVM

```
GET "/api/svm/peer-permissions/?svm.name=VS1"
```

1. Retrieves a list of SVM peer permissions of a specific cluster peer

```
GET "/api/svm/peer-permissions/?cluster_peer.name=cluster2"
```

#### Learn more

- DOC /svm/peer-permissions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster_peer.name | string | query | False | Filter by cluster_peer.name |
| cluster_peer.uuid | string | query | False | Filter by cluster_peer.uuid |
| applications | string | query | False | Filter by applications |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[svm_peer_permission] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "applications": [
        "snapmirror",
        "flexcache"
      ],
      "cluster_peer": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster2",
        "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345574 | Failed to find the SVM or volume name with UUID. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster_peer

Peer cluster details

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*".

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer_permission

Manage SVM peer permissions.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| cluster_peer | cluster_peer | Peer cluster details |
| svm | svm | Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*". |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an SVM peer permission

POST `/svm/peer-permissions`

**Introduced In:** 9.6

Creates an SVM peer permission.

#### Required properties

- `svm.name` or `svm.uuid` - SVM name
- `cluster_peer.uuid` or `cluster_peer.name` - Peer cluster name or peer cluster UUID

- `applications` - Peering applications

#### Related ONTAP commands

- `vserver peer permission create`

#### Examples

The following examples show how to create SVM peer permissions.

1. Creates an SVM peer permission entry with the local SVM and cluster peer names

```
 POST "/api/svm/peer-permissions" '{"cluster_peer":{"name":"cluster2"},
"svm":{"name":"VS1"}, "applications":["snapmirror"]}'
```

1. Creates an SVM peer permission entry with the local SVM and cluster peer UUID

```
 POST "/api/svm/peer-permissions" '{"cluster_peer":{"uuid":"d3268a74-ee76-
11e8-a9bb-005056ac6dc9"}, "svm":{"uuid":"8f467b93-f2f1-11e8-9027-
005056ac81fc"}, "applications":["snapmirror"]}'
```

1. Creates an SVM peer permission entry with all SVMs and the cluster peer name

```
 POST "/api/svm/peer-permissions" '{"cluster_peer":{"name":"cluster2"},
"svm":{"name":"*"}, "applications":["snapmirror"]}'
```

#### Learn more

- DOC /svm/peer-permissions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |

| Name | Type | Description |
| --- | --- | --- |
| cluster_peer | cluster_peer | Peer cluster details |
| svm | svm | Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*". |

#### Example request

```
{
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| cluster_peer | cluster_peer | Peer cluster details |
| svm | svm | Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*". |

### Example response

```
{
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345572 | {field} is a required field. |
| 26345573 | Failed to find the SVM or volume UUID with name. |
| 26345574 | Failed to find the SVM or volume name with UUID. |
| 26345575 | The specified peer cluster name and peer cluster UUID do not match. |
| 9896057 | SVM peer permission already exists. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster_peer

Peer cluster details

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*".

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer_permission

Manage SVM peer permissions.

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |
| cluster_peer | cluster_peer | Peer cluster details |
| svm | svm | Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*". |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete SVM peer permissions

DELETE `/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`

**Introduced In:** 9.6

Deletes the SVM peer permissions.

#### Related ONTAP commands

- `vserver peer permission delete`

#### Example

Deletes an SVM peer permission.

```
DELETE "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-
005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
```

#### Learn more

- DOC /svm/peer-permissions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster_peer.uuid | string | path | True | Peer cluster UUID<br>• Introduced in: 9.7 |
| svm.uuid | string | path | True | SVM UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345574 | Failed to find the SVM or volume name with UUID. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an SVM peer permission instance

GET `/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`

**Introduced In:** 9.6

Retrieves the SVM peer permission instance.

#### Related ONTAP commands

- `vserver peer permission show`

#### Example

The following example shows how to retrieve the parameters for an SVM peer permission.

```
GET "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-
005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
```

#### Learn more

- DOC /svm/peer-permissions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster_peer.uuid | string | path | True | Peer cluster UUID<br>• Introduced in: 9.7 |
| svm.uuid | string | path | True | SVM UUID |
| fields | array[string] | query | False | Specify the fields to return. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| cluster_peer | cluster_peer | Peer cluster details |
| svm | svm | Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*". |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345574 | Failed to find the SVM or volume name with UUID. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster_peer

Peer cluster details

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*".

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update SVM peer permissions

PATCH `/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`

**Introduced In:** 9.6

Updates the SVM peer permissions.

#### Related ONTAP commands

- `vserver peer permission modify`

#### Example

Updates an SVM peer permission.

```
PATCH "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-
005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
'{"applications":["flexcache"]}'
```

#### Learn more

- DOC /svm/peer-permissions

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cluster_peer.uuid | string | path | True | Peer cluster UUID<br>• Introduced in: 9.7 |
| svm.uuid | string | path | True | SVM UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |

#### Example request

```
{
  "applications": [
    "snapmirror",
    "flexcache"
  ]
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |

#### Example response

```
{
  "applications": [
    "snapmirror",
    "flexcache"
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
| --- | --- |
| 26345572 | {field} is a required field. |
| 26345574 | Failed to find the SVM or volume name with UUID. |
| 9896059 | SVM peer permission does not exist. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster_peer

Peer cluster details

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

Local SVM permitted for peer relationship. To create peer permissions for all SVMs, specify the SVM name as "*".

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer_permission

Manage SVM peer permissions.

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SVM peer relationships

### Manage SVM peer relationships

#### Managing SVM peers

The SVM peer commands allow you to create and manage SVM peering relationships.

### SVM peer APIs

The following APIs are used to manage SVM peers:

- GET /api/svm/peers
- POST /api/svm/peers
- GET /api/svm/peers/{uuid}
- PATCH /api/svm/peers/{uuid}
- DELETE /api/svm/peers/{uuid}

### Retrieve SVM peer relationships

GET `/svm/peers`

**Introduced In:** 9.6

Retrieves the list of SVM peer relationships.

#### Related ONTAP commands

- `vserver peer show`

#### Examples

The following examples show how to retrieve a collection of SVM peer relationships based on a query.

1. Retrieves a list of SVM peers of a specific local SVM

```
GET "/api/svm/peers/?svm.name=VS1"
```

1. Retrieves a list of SVM peers of a specific cluster peer

```
GET "/api/svm/peers/?peer.cluster.name=cluster2"
```

**Learn more**

- DOC /svm/peers

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| state | string | query | False | Filter by state |
| name | string | query | False | Filter by name |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| applications | string | query | False | Filter by applications |
| peer.cluster.name | string | query | False | Filter by peer.cluster.name |
| peer.cluster.uuid | string | query | False | Filter by peer.cluster.uuid |
| peer.svm.name | string | query | False | Filter by peer.svm.name |
| peer.svm.uuid | string | query | False | Filter by peer.svm.uuid |
| uuid | string | query | False | Filter by uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[svm_peer] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "applications": [
        "snapmirror",
        "lun_copy"
      ],
      "force": true,
      "name": "string",
      "peer": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster2",
          "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "state": "peered",
```

```
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string"
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345578 | Internal error. Unable to retrieve local or peer SVM name. |
| 9896086 | Peer SVM name conflicts with one of the following: a peer SVM in an existing SVM peer relationship, a local SVM, or an IPSpace. Use the "name" property to uniquely specify the peer SVM alias name. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

peer

Details for a peer SVM object.

| Name | Type | Description |
| --- | --- | --- |
| cluster | cluster | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

svm

Local SVM details

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer

An SVM peer relationship object.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |
| peer | peer | Details for a peer SVM object. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | SVM peering state. To accept a pending SVM peer request, PATCH the state to "peered". To reject a pending SVM peer request, PATCH the state to "rejected". To suspend a peered SVM peer relationship, PATCH the state to "suspended". To resume a suspended SVM peer relationship, PATCH the state to "peered". The states "initiated", "pending", and "initializing" are system-generated and cannot be used for PATCH. |
| svm | svm | Local SVM details |
| uuid | string | SVM peer relationship UUID |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new SVM peer relationship**

POST `/svm/peers`

**Introduced In:** 9.6

Creates a new SVM peer relationship.

#### Important notes

- The create request accepts peer SVM name as input instead of peer SVM UUID as the local cluster cannot validate peer SVM based on UUID.
- The input parameter `name` refers to the local name of the peer SVM. The `peer cluster name` parameter is optional for creating intracluster SVM peer relationships.

#### Required properties

- `svm.name` or `svm.uuid` - SVM name or SVM UUID
- `peer.svm.name` or `peer.svm.uuid` - Peer SVM name or Peer SVM UUID
- `peer.cluster.name` or `peer.cluster.uuid` - Peer cluster name or peer cluster UUID
- `applications` - Peering applications

#### Related ONTAP commands

- `vserver peer create`

#### Example

Creates a new SVM peer relationship.

```
POST "/api/svm/peers" '{"svm":{"name":"vs1"},
"peer.cluster.name":"cluster2", "peer.svm.name":"VS1",
"applications":["snapmirror"]}'
```

#### Learn more

- DOC /svm/peers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |

| Name | Type | Description |
| --- | --- | --- |
| force | boolean | Use this to suspend, resume or delete the SVM peer relationship even if the remote cluster is not accessible due to, for example, network connectivity issues. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |
| peer | peer | Details for a peer SVM object. |
| svm | svm | Local SVM details |
| uuid | string | SVM peer relationship UUID |

**Example request**

```
{
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "name": "string",
  "peer": {
    "cluster": {
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 13434889 | Internal error. Wait and retry. |
| 26345575 | The specified peer cluster name and peer cluster UUID do not match. |
| 26345579 | The specified field is invalid. |
| 26345580 | SVM name or SVM UUID must be provided. |
| 9896086 | Peer SVM name conflicts with one of the following: a peer SVM in an existing SVM peer relationship, a local SVM, or an IPSpace. Use the "name" property to uniquely specify the peer SVM alias name. |

| Error codes | Description |
| --- | --- |
| 9895996 | Cannot specify lun-copy as an inter-cluster application. |
| 9896093 | The flexcache application is not supported on this platform. |
| 9896095 | SVM name is in use by a peer Vserver in a Vserver peer relationship. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

peer

Details for a peer SVM object.

| Name | Type | Description |
|---|---|---|
| cluster | cluster | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

svm

Local SVM details

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer

An SVM peer relationship object.

| Name | Type | Description |
| --- | --- | --- |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| force | boolean | Use this to suspend, resume or delete the SVM peer relationship even if the remote cluster is not accessible due to, for example, network connectivity issues. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |
| peer | peer | Details for a peer SVM object. |
| svm | svm | Local SVM details |
| uuid | string | SVM peer relationship UUID |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an SVM peer relationship

DELETE `/svm/peers/{uuid}`

**Introduced In:** 9.6

Deletes the SVM peer relationship.

#### Related ONTAP commands

- `vserver peer delete`

#### Example

1. Deletes an SVM peer relationship.

```
DELETE "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
```

1. Deletes an SVM peer relationship using force flag

```
DELETE "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9" '{"force":
"true"}'
```

#### Learn more

- DOC /svm/peers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | SVM peer relationship UUID<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Use this parameter to delete the SVM peer relationship even when the remote cluster is not accessible. For example, if there are network connectivity issues.<br>• Introduced in: 9.17 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345578 | Internal error. Unable to retrieve local or peer SVM name. |
| 9895956 | Cannot delete an SVM that is part of an SVM peer or transition peer relationship. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve an SVM peer relationship instance**

GET `/svm/peers/{uuid}`

**Introduced In:** 9.6

Retrieves the SVM peer relationship instance.

#### Related ONTAP commands

- `vserver peer show`

#### Example

Retrieves the parameters of an SVM peer relationship.

```
GET "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
```

#### Learn more

- DOC /svm/peers

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | SVM peer relationship UUID<br>• Introduced in: 9.7 |
| fields | array[string] | query | False | Specify the fields to return. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |
| peer | peer | Details for a peer SVM object. |
| state | string | SVM peering state. To accept a pending SVM peer request, PATCH the state to "peered". To reject a pending SVM peer request, PATCH the state to "rejected". To suspend a peered SVM peer relationship, PATCH the state to "suspended". To resume a suspended SVM peer relationship, PATCH the state to "peered". The states "initiated", "pending", and "initializing" are system-generated and cannot be used for PATCH. |
| svm | svm | Local SVM details |
| uuid | string | SVM peer relationship UUID |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "force": true,
  "name": "string",
  "peer": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "peered",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 26345578 | Internal error. Unable to retrieve local or peer SVM name. |
| 9896086 | Peer SVM name conflicts with one of the following: a peer SVM in an existing SVM peer relationship, a local SVM, or an IPSpace. Use the "name" property to uniquely specify the peer SVM alias name. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

cluster

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

peer

Details for a peer SVM object.

| Name | Type | Description |
|---|---|---|
| cluster | cluster | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

svm

Local SVM details

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an SVM peer relationship

PATCH `/svm/peers/{uuid}`

**Introduced In:** 9.6

Updates the SVM peer relationship.

#### Related ONTAP commands

- `vserver peer modify`

#### Examples

The following examples show how to update an SVM peer relationship. The input parameter 'name' refers to the local name of the peer SVM.

1. Accepts an SVM peer relationship

```
 PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"state":"peered"}'
```

1. Updates the local name of an SVM peer relationship

```
 PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"name":"vs2"}'
```

1. Suspends an SVM peer relationship using force flag

```
 PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"state":"suspended", "force": "true"}'
```

**Learn more**

- DOC /svm/peers

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | SVM peer relationship UUID<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| applications | array[string] | A list of applications for an SVM peer relationship. |
| force | boolean | Use this to suspend, resume or delete the SVM peer relationship even if the remote cluster is not accessible due to, for example, network connectivity issues. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | SVM peering state. To accept a pending SVM peer request, PATCH the state to "peered". To reject a pending SVM peer request, PATCH the state to "rejected". To suspend a peered SVM peer relationship, PATCH the state to "suspended". To resume a suspended SVM peer relationship, PATCH the state to "peered". The states "initiated", "pending", and "initializing" are system-generated and cannot be used for PATCH. |
| uuid | string | SVM peer relationship UUID |

**Example request**

```
{
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "name": "string",
  "state": "peered",
  "uuid": "string"
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 13434889 | Internal error. Wait and retry. |
| 26345575 | The specified peer cluster name and peer cluster UUID do not match. |
| 26345576 | Given peer state is invalid. |
| 26345577 | One of the following is required: applications, state, or name. |
| 26345578 | Internal error. Unable to retrieve local or peer SVM name. |
| 26345579 | The specified field is invalid. |
| 26345581 | Peer cluster name could not be retrieved or validated. |
| 9896077 | The peer relationship is in use by FlexCache. View the FlexCache relationships, delete them and retry the operation. |
| 9896088 | System generated a name for the peer SVM because of a naming conflict. Use the name property to uniquely identify the peer SVM alias name. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

cluster

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

peer

Details for a peer SVM object.

| Name | Type | Description |
|---|---|---|
| cluster | cluster | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

svm

Local SVM details

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

svm_peer

An SVM peer relationship object.

| Name | Type | Description |
| --- | --- | --- |
| applications | array[string] | A list of applications for an SVM peer relationship. |
| force | boolean | Use this to suspend, resume or delete the SVM peer relationship even if the remote cluster is not accessible due to, for example, network connectivity issues. |
| name | string | A peer SVM alias name to avoid a name conflict on the local cluster. |
| state | string | SVM peering state. To accept a pending SVM peer request, PATCH the state to "peered". To reject a pending SVM peer request, PATCH the state to "rejected". To suspend a peered SVM peer relationship, PATCH the state to "suspended". To resume a suspended SVM peer relationship, PATCH the state to "peered". The states "initiated", "pending", and "initializing" are system-generated and cannot be used for PATCH. |
| uuid | string | SVM peer relationship UUID |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage SVMs

### Manage SVMs

#### Managing SVMs

Cluster administrators can manage any SVM bound to the cluster. In addition, SVMs can also be managed by their SVM administrators. The SVM administrator manages the SVM resources, such as volumes, protocols and services, depending on the capabilities assigned by the cluster administrator. SVM administrators cannot create, modify, or delete SVMs. The cluster administrator manages SVM create, modify, or delete operations.

While configuring CIFS, you must also configure IP interfaces and DNS. No other protocol configuration is allowed when configuring NVMe. NFS, FCP, CIFS, iSCSI, and S3 protocols can be configured together.

SVM administrators might have all or some of the following administration capabilities:

1. Data access protocol configuration
   Configures data access protocols, such as NFS, CIFS, iSCSI, S3, and Fibre Channel (FC) protocols (Fibre Channel over Ethernet included).
2. Services configuration
   Configures services such as LDAP, NIS, and DNS.
3. Monitoring SVM
   Monitors jobs, network connections, network interfaces, and SVM health.
4. Updating the TLS certificate for this SVM.

### Retrieve SVMs and SVM properties

GET `/svm/svms`

**Introduced In:** 9.6

Retrieves a list of SVMs and individual SVM properties. This includes protocol configurations such as CIFS, NFS and S3, export policies, name service configurations, and network services.

#### Important notes

- The SVM object includes a large set of fields and can be expensive to retrieve. Use this API to list the collection of SVMs, and to retrieve only the full details of individual SVMs as needed.
- It is not recommended to create or delete more than five SVMs in parallel.
- REST APIs only expose a data SVM as an SVM.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `snapmirror.*`

#### Related ONTAP commands

- `vserver show`

#### Examples

- Retrieves a list of SVMs in the cluster sorted by name.

```
GET "/api/svm/svms?order_by=name"
```

- Retrieves a list of SVMs in the cluster that have the NFS protocol enabled.

```
GET "/api/svm/svms?nfs.enabled=true"
```

- Retrieves a list of SVMs in the cluster that have the CIFS protocol enabled.

```
GET "/api/svm/svms?cifs.enabled=true"
```

- Retrieves a list of SVMs in the cluster that have the S3 protocol enabled.

```
GET "/api/svm/svms?s3.enabled=true"
```

- Retrieves a list of SVMs in the cluster that have the FCP protocol allowed.

```
GET "/api/svm/svms?fcp.allowed=true"
```

- Retrieves a list of SVMs in the cluster that have the CIFS protocol allowed.

```
GET "/api/svm/svms?cifs.allowed=true"
```

- Retrieves a list of SVMs in the cluster where the NDMP protocol is specified as allowed.

```
GET "/api/svm/svms?ndmp.allowed=true"
```

- Retrieves a list of SVMs in the cluster that have the s3 protocol allowed.

```
GET "/api/svm/svms?s3.allowed=true"
```

**Learn more**

- DOC /svm/svms

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nis.enabled | boolean | query | False | Filter by nis.enabled |
| nis.servers | string | query | False | Filter by nis.servers |
| nis.domain | string | query | False | Filter by nis.domain |
| nvme.allowed | boolean | query | False | Filter by nvme.allowed<br>• Introduced in: 9.9 |
| nvme.enabled | boolean | query | False | Filter by nvme.enabled |
| language | string | query | False | Filter by language |
| nfs.allowed | boolean | query | False | Filter by nfs.allowed<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nfs.enabled | boolean | query | False | Filter by nfs.enabled |
| comment | string | query | False | Filter by comment |
| aggregates.name | string | query | False | Filter by aggregates.name |
| aggregates.uuid | string | query | False | Filter by aggregates.uuid |
| subtype | string | query | False | Filter by subtype |
| dns.servers | string | query | False | Filter by dns.servers |
| dns.domains | string | query | False | Filter by dns.domains |
| fcp.allowed | boolean | query | False | Filter by fcp.allowed<br>• Introduced in: 9.9 |
| fcp.enabled | boolean | query | False | Filter by fcp.enabled |
| iscsi.allowed | boolean | query | False | Filter by iscsi.allowed<br>• Introduced in: 9.9 |
| iscsi.enabled | boolean | query | False | Filter by iscsi.enabled |
| name | string | query | False | Filter by name |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid |
| ipspace.name | string | query | False | Filter by ipspace.name |
| ldap.base_dn | string | query | False | Filter by ldap.base_dn |
| ldap.servers | string | query | False | Filter by ldap.servers |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ldap.enabled | boolean | query | False | Filter by ldap.enabled |
| ldap.bind_dn | string | query | False | Filter by ldap.bind_dn |
| ldap.ad_domain | string | query | False | Filter by ldap.ad_domain |
| uuid | string | query | False | Filter by uuid |
| cifs.name | string | query | False | Filter by cifs.name |
| cifs.ad_domain.fqdn | string | query | False | Filter by cifs.ad_domain.fqdn |
| cifs.ad_domain.organizational_unit | string | query | False | Filter by cifs.ad_domain.organizational_unit |
| cifs.allowed | boolean | query | False | Filter by cifs.allowed<br>• Introduced in: 9.9 |
| cifs.enabled | boolean | query | False | Filter by cifs.enabled |
| s3.name | string | query | False | Filter by s3.name<br>• Introduced in: 9.7 |
| s3.allowed | boolean | query | False | Filter by s3.allowed<br>• Introduced in: 9.12 |
| s3.enabled | boolean | query | False | Filter by s3.enabled<br>• Introduced in: 9.7 |
| s3.is_http_enabled | boolean | query | False | Filter by s3.is_http_enabled<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| s3.is_https_enabled | boolean | query | False | Filter by s3.is_https_enabled<br>• Introduced in: 9.16 |
| s3.port | integer | query | False | Filter by s3.port<br>• Introduced in: 9.16 |
| s3.secure_port | integer | query | False | Filter by s3.secure_port<br>• Introduced in: 9.16 |
| s3.certificate.uuid | string | query | False | Filter by s3.certificate.uuid<br>• Introduced in: 9.16 |
| s3.certificate.name | string | query | False | Filter by s3.certificate.name<br>• Introduced in: 9.16 |
| ndmp.allowed | boolean | query | False | Filter by ndmp.allowed<br>• Introduced in: 9.10 |
| nsswitch.netgroup | string | query | False | Filter by nsswitch.netgroup |
| nsswitch.group | string | query | False | Filter by nsswitch.group |
| nsswitch.hosts | string | query | False | Filter by nsswitch.hosts |
| nsswitch.namemap | string | query | False | Filter by nsswitch.namemap |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nsswitch.passwd | string | query | False | Filter by nsswitch.passwd |
| qos_policy.name | string | query | False | Filter qos_policy.name<br>• Introduced in: 9.9 |
| qos_policy.uuid | string | query | False | Filter qos_policy.uuid<br>• Introduced in: 9.9 |
| max_volumes | string | query | False | Filter max_volumes<br>• Introduced in: 9.9 |
| fc_interfaces.name | string | query | False | Filter by fc_interfaces.name<br>• Introduced in: 9.7 |
| fc_interfaces.data_protocol | string | query | False | Filter by fc_interfaces.data_protocol<br>• Introduced in: 9.7 |
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid |
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name |
| state | string | query | False | Filter by state |
| certificate.uuid | string | query | False | Filter by certificate.uuid<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware_default_volume_state | string | query | False | Filter by anti_ransomware_default_volume_state<br>• Introduced in: 9.10 |
| qos_adaptive_policy_group_template.name | string | query | False | Filter by qos_adaptive_policy_group_template.name<br>• Introduced in: 9.13 |
| qos_adaptive_policy_group_template.uuid | string | query | False | Filter by qos_adaptive_policy_group_template.uuid<br>• Introduced in: 9.13 |
| qos_policy_group_template.name | string | query | False | Filter by qos_policy_group_template.name<br>• Introduced in: 9.16 |
| qos_policy_group_template.uuid | string | query | False | Filter by qos_policy_group_template.uuid<br>• Introduced in: 9.16 |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | query | False | Filter by anti_ransomware_auto_switch_from_learning_to_enabled<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware_auto_switch_minimum_incoming_data | string | query | False | Filter by anti_ransomware_auto_switch_minimum_incoming_data<br><br>• Introduced in: 9.13 |
| anti_ransomware_auto_switch_minimum_incoming_data_percent | integer | query | False | Filter by anti_ransomware_auto_switch_minimum_incoming_data_percent<br><br>• Introduced in: 9.13 |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | query | False | Filter by anti_ransomware_auto_switch_duration_without_new_file_extension<br><br>• Introduced in: 9.13 |
| anti_ransomware_auto_switch_minimum_learning_period | integer | query | False | Filter by anti_ransomware_auto_switch_minimum_learning_period<br><br>• Introduced in: 9.13 |
| anti_ransomware_auto_switch_minimum_file_count | integer | query | False | Filter by anti_ransomware_auto_switch_minimum_file_count<br><br>• Introduced in: 9.13 |
| anti_ransomware_auto_switch_minimum_file_extension | integer | query | False | Filter by anti_ransomware_auto_switch_minimum_file_extension<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| auto_enable_analytics | boolean | query | False | Filter by auto_enable_analytics.<br>• Introduced in: 9.12 |
| auto_enable_activity_tracking | boolean | query | False | Filter by auto_enable_activity_tracking.<br>• Introduced in: 9.12 |
| storage.allocated | integer | query | False | Filter by storage_allocated<br>• Introduced in: 9.14 |
| storage.available | integer | query | False | Filter by storage_available<br>• Introduced in: 9.14 |
| storage.used_percentage | integer | query | False | Filter by storage_used_percentage<br>• Introduced in: 9.14 |
| number_of_volumes_in_recovery_queue | integer | query | False | Filter by number_of_volumes_in_recovery_queue<br>• Introduced in: 9.13 |
| total_volume_size_in_recovery_queue | integer | query | False | Filter by total_volume_size_in_recovery_queue<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| storage.limit_threshold_exceeded | integer | query | False | Filter by storage.limit_threshold_exceeded<br><br>• Introduced in: 9.14 |
| storage.limit | integer | query | False | Filter by storage.limit<br><br>• Introduced in: 9.14 |
| storage.limit_threshold_alert | integer | query | False | Filter by storage.limit_threshold_alert<br><br>• Introduced in: 9.14 |
| snapshot_autodelete_enabled | boolean | query | False | Filter by snapshot_autodelete_enabled<br><br>• Introduced in: 9.18 |
| snapshot_reserve_percent | integer | query | False | Filter by snapshot_reserve_percent<br><br>• Introduced in: 9.18 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[svm] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "available_size": 10156560384,
          "name": "aggr1",
          "snaplock_type": "string",
          "state": "string",
          "type": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "anti_ransomware_auto_switch_minimum_incoming_data": "string",
      "anti_ransomware_default_volume_state": "string",
      "anti_ransomware_incoming_write_threshold": "string",
      "anti_ransomware_incoming_write_threshold_percent": "string",
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "cifs": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ad_domain": {
          "default_site": "string",
          "fqdn": "example.com",
          "organizational_unit": "string"
        },
        "auth-style": "domain",
        "domain_workgroup": "string",
        "name": "CIFS1",
        "workgroup": "workgrp1"
      },
      "comment": "string",
      "dns": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "domains": [
          "example.com",
          "example2.example3.com"
        ],
        "servers": [
          "10.224.65.20",
          "2001:db08:a0b:12f0::1"
        ]
      },
      "fc_interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "data_protocol": "string",
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "fcp": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "ip_interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
          "services": [
            "data_nfs"
          ],
          "subnet": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "subnet1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "iscsi": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        }
      },
      "language": "c.utf_8",
      "ldap": {
        "ad_domain": "string",
        "base_dn": "string",
        "bind_dn": "string",
        "servers": [
          "string"
        ]
      },
      "max_volumes": "string",
      "name": "svm1",
      "nfs": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "nis": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "domain": "string",
        "servers": [
          "string"
        ]
      },
      "nsswitch": {
        "group": [
          "string"
        ],
        "hosts": [
          "string"
        ],
        "namemap": [
          "string"
        ],
        "netgroup": [
          "string"
        ],
        "passwd": [
          "string"
```

```
        ]
      },
      "number_of_volumes_in_recovery_queue": 0,
      "nvme": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "qos_adaptive_policy_group_template": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "qos_policy_group_template": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
```

```
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "s3": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "certificate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "default_unix_user": "string",
        "default_win_user": "string",
        "name": "s3-server-1"
      },
      "snapmirror": {
        "protected_consistency_group_count": 0,
        "protected_volumes_count": 0
      },
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "running",
      "storage": {
        "allocated": 0,
```

```
          "available": 0,
          "used_percentage": 0
        },
        "subtype": "string",
        "total_volume_size_in_recovery_queue": 0,
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    ]
  }
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| available_size | integer | Space available, in bytes. |
| name | string | |
| snaplock_type | string | SnapLock type. |
| state | string | Aggregate state. |
| type | string | Type of aggregate. |
| uuid | string | |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

arw_vserver

Anti-ransomware related information for the SVM.

| Name | Type | Description |
|---|---|---|
| event_log | event_log | |

certificate

Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. |

cifs

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ad_domain | ad_domain | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the CIFS service. If it is false, only the cluster administrator can manage the service. |

| Name | Type | Description |
| --- | --- | --- |
| auth-style | string | Authentication type. |
| domain_workgroup | string | The NetBIOS name of the domain or workgroup associated with the CIFS server. |
| enabled | boolean | If allowed, setting to true enables the CIFS service. |
| name | string | The NetBIOS name of the CIFS server. |
| workgroup | string | The workgroup name. |

dns

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |

node

The node on which the FC port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the FC port is located. |

fc_port_reference

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the Fibre Channel interface is defined by the location of its port.

| Name | Type | Description |
| --- | --- | --- |
| port | fc_port_reference | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

fc_interface_svm

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| data_protocol | string | The data protocol for which the Fibre Channel interface is configured. |
| name | string | The name of the Fibre Channel interface. |
| uuid | string | The unique identifier of the Fibre Channel interface. |

fcp

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the FCP service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the FCP service. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

port_svm

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

location

Home_node is optional.

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. |
| home_node | home_node | |
| home_port | port_svm | |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

ip_interface_svm

Interface parameters. Name and home_node are optional.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface (optional). |
| services | array[string] | The services associated with the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Either the UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

iscsi

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the iSCSI service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the iSCSI service. |

ldap

| Name | Type | Description |
|---|---|---|
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST. |
| base_dn | string | Specifies the default base DN for all searches. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. SVM API supports anonymous binding. For Simple and SASL LDAP binding, use the LDAP API endpoint. |
| enabled | boolean | Enable LDAP? Setting to true creates a configuration if not already created. |

| Name | Type | Description |
| --- | --- | --- |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| servers | array[string] | |

ndmp

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NDMP service. If it is false, only the cluster administrator can manage the service. |

nfs

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NFS service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the NFS service. |

nis

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| domain | string | The NIS domain to which this configuration belongs. |
| enabled | boolean | Enable NIS? Setting to true creates a configuration if not already created. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

nsswitch

Name service switch configuration

| Name | Type | Description |
|---|---|---|
| group | array[string] | Group sources |
| hosts | array[string] | Host sources |
| namemap | array[string] | NameMap sources |
| netgroup | array[string] | NetGroup sources |
| passwd | array[string] | Password sources |

nvme

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NVMe service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the NVMe service. |

qos_adaptive_policy_group_template

This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space.

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy

This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy_group_template

This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs).

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

network_route_for_svm

| Name | Type | Description |
| --- | --- | --- |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

s3

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether or not to enable S3. Setting this value to true creates a service if one is not yet created. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server. |

| Name | Type | Description |
| --- | --- | --- |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server. |
| name | string | Specifies the name of the S3 server. A server name length can range from 1 to 253 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. |

snapmirror

Specifies attributes for SVM DR protection.

| Name | Type | Description |
| --- | --- | --- |
| is_protected | boolean | Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data. |
| protected_consistency_group_count | integer | Specifies the number of SVM DR protected consistency groups in the SVM. |
| protected_volumes_count | integer | Specifies the number of SVM DR protected volumes in the SVM. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

storage

| Name | Type | Description |
| --- | --- | --- |
| allocated | integer | Total size of the volumes in SVM, in bytes. |
| available | integer | Currently available storage capacity in SVM, in bytes. |
| limit | integer | Maximum storage permitted on a single SVM, in bytes. |
| limit_threshold_alert | integer | Indicates at what percentage of storage capacity an alert message is sent. The default value is 90. |
| limit_threshold_exceeded | boolean | Indicates whether the total storage capacity exceeds the alert percentage. |
| used_percentage | integer | The percentage of storage capacity used. |

svm

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests. |
| cifs | cifs | |
| comment | string | Comment |
| dns | dns | |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| ipspace | ipspace | Either the UUID or name may be supplied on input. |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| ldap | ldap | |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |
| nis | nis | |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |
| qos_policy | qos_policy | This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| s3 | s3 | |

| Name | Type | Description |
| --- | --- | --- |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| state | string | SVM State |
| storage | storage | |
| subtype | string | SVM subtype. The SVM subtype sync_destination is created automatically when an SVM of subtype sync_source is created on the source MetroCluster cluster. A POST request with sync_destination as SVM subtype is invalid. SVM of subtype data_engine cannot be explicitly created by the admin and most management changes are not allowed on it. |
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create and provision an SVM

POST `/svm/svms`

**Introduced In:** 9.6

Creates and provisions an SVM. If no IPspace is provided, then the SVM is created on the `Default` IPspace.

- The number of parallel SVMs that can be created must not be greater than five.
- If a sixth SVM POST request is issued, the following error message is generated: "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again."

#### Required properties

- `name` - Name of the SVM to be created.

#### Recommended optional properties

- `ipspace.name` or `ipspace.uuid` - IPspace of the SVM
- `is_space_reporting_logical` - Logical Space Reporting parameter of the SVM
- `is_space_enforcement_logical` - Logical Space Enforcement parameter of the SVM
- `ip_interfaces` - If provided, the following fields are required:
- `ip_interfaces.name` - Name of the interface
- `ip_interfaces.ip.address` - IP address
- `ip_interfaces.ip.netmask` - Netmask length or IP address
- `ip_interfaces.location.broadcast_domain.uuid` or `ip_interfaces.location.broadcast_domain.name` - Broadcast domain name or UUID belonging to the same IPspace of the SVM.
- `ip_interfaces.location.home_port.name` - Home port name
- `ip_interfaces.location.home_port.uuid` - Home port uuid
- `subnet.uuid` or `subnet.name` - Either name or UUID of the subnet to create.
- `routes` - If provided, the following field is required:
  - `routes.gateway` - Gateway IP address
- `cifs` - If provided, interfaces, routes and DNS must be provided. The following fields are also required:
  - `cifs.name` - Name of the CIFS server to be created for the SVM.
  - `cifs.ad_domain.fqdn` - Fully qualified domain name
  - `cifs.ad_domain.user` - Administrator username

  - `cifs.ad_domain.password` - User password
- `ldap` - If provided, the following fields are required:
  - `ldap.servers` or `ldap.ad_domain` - LDAP server list or Active Directory domain
  - `ldap.bind_dn` - Bind DN
  - `ldap.base_dn` - Base DN
- `nis` - If provided, the following fields are required:
  - `nis.servers` - NIS servers
  - `nis.domain` - NIS domain
- `dns` - If provided, the following fields are required:
  - `dns.servers` - Name servers
  - `dns.domains` - Domains
- `fc_interfaces` - If provided, the following fields are required:
  - `fc_interfaces.name` - Fibre Channel interface name
  - `fc_interfaces.data_protocol` - Fibre Channel interface data protocol
  - `fc_interfaces.location.port.uuid` or `fc_interfaces.location.port.name` and `fc_interfaces.location.port.node.name` - Either port UUID or port name and node name together must be provided.
- `s3` - If provided, the following field should also be specified:
  - <code>s3.name</code> - Name of the S3 server. If <code>s3.name' is not specified while </code>s3.enabled` is set to 'true', the S3 server will be created with the default name '<svm.name>_S3Server'.</svm.name>
  - `s3.port` - S3 server listener port.
  - `s3.secure_port` - S3 server listener port for HTTPS.
  - `s3.is_http_enabled` - S3 server connections over HTTP.
  - `s3.is_https_enabled` - S3 server connections over HTTPS.
  - `s3.certificate.name` - S3 server certificate name. This is required if the S3 server runs over HTTPS.
  - `s3.certificate.uuid` - S3 server certificate UUID. This is required if the S3 server runs over HTTPS.
- `auto_enable_analytics` - Auto-enable file system analytics on new volumes created in the SVM.
- `auto_enable_activity_tracking` - Auto-enable volume activity-tracking on new volumes created in the SVM.
- `storage.limit` - Maximum storage permitted on a single SVM.
- `storage.limit_threshold_alert` - At what percentage of storage capacity, alert message needs to be sent.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `language` - *C.UTF-8*
- `ipspace.name` - *Default*
- `snapshot_policy.name` - *Default*
- `subtype` - *Default* ( *sync-source* if MetroCluster configuration )
- `anti_ransomware_default_volume_state` - *disabled*
- `qos_adaptive_policy_group_template` - *extreme* ( if using a platform with disaggregated storage and neither qos_policy_group_template nor qos_adaptive_policy_group_template are provided)

#### Related ONTAP commands

- `vserver create`
- `vserver add-aggregates`
- `network interface create`
- `network route create`
- `vserver services name-service dns create`
- `vserver nfs create`
- `vserver services name-service ldap client create`
- `vserver cifs create`
- `vserver services name-service nis-domain create`
- `vserver iscsi create`
- `vserver nvme create`
- `vserver fcp create`
- `vserver services name-service ns-switch create`
- `vserver object-store-server create`
- `vserver add-protocols`
- `vserver remove-protocols`

#### Examples

- Creates an SVM with default "snapshot_policy".

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"}}'
```

- Creates an SVM and configures NFS, CIFS, and S3.

```
POST "/api/svm/svms" '{"name":"testVs", "nfs":{"enabled":"true"},
"cifs":{"enabled":"true"}, "s3":{"enabled":"true"}}'
```

- Creates an SVM and configures NVMe.

```
POST "/api/svm/svms" '{"name":"testVs", "nvme":{"enabled":"true"}}'
```

- Creates an SVM and configures LDAP.

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"},
"ldap":{"servers":["10.140.101.1","10.140.101.2"],
"ad_domain":"abc.com", "base_dn":"dc=netapp,dc=com",
"bind_dn":"dc=netapp,dc=com"}}'
```

- Creates an SVM and configures NIS.

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"}, "nis":{"enabled":"true",
"domain":"def.com","servers":["10.224.223.130", "10.224.223.131"]}}'
```

- Creates an SVM and configures DNS.

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"},
"dns":{"domains":["abc.com","def.com"], "servers":["10.224.223.130",
"10.224.223.131"]}}'
```

- Creates an SVM and configures a LIF.

```
POST "/api/svm/svms" '{"name":"testVs", "ip_interfaces":
[{"name":"lif1", "ip":{"address":"10.10.10.7", "netmask":
"255.255.255.0"}, "location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-
management"}]}'
```

- Creates an SVM and configures a LIF with IPV6 address.

```
POST "/api/svm/svms" '{"name":"testVs", "ip_interfaces":
[{"name":"lif2",
"ip":{"address":"fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b",
"netmask":"24"}, "location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-
management"}]}'
```

- Creates an SVM and configures CIFS.

```
POST "/api/svm/svms" '{"name":"testVs", "cifs":{"name":"CIFDOC",
"ad_domain":{"fqdn":"abc.def.com", "organizational_unit":"CN=Computers",
"user":"cif_admin", "password":"abc123"}},
"ip_interfaces":[{"name":"lif1", "ip":{"address":"10.10.10.7",
"netmask": "255.255.255.0"},
"location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-
management"}],"routes": [{"destination": {"address": "0.0.0.0",
"netmask": "0"}, "gateway": "10.10.10.7"}],
"dns":{"domains":["abc.def.com", "def.com"],
"servers":["10.224.223.130", "10.224.223.131"]}}'
```

- Creates an SVM with an S3 server enabled and configured.

```
POST "/api/svm/svms" '{"name":"svm5", "s3":{"name":"s3-server-1",
"enabled":true, "allowed":true, "is_http_enabled": true,
"is_https_enabled":false}}'
```

- Creates an SVM and disallows NVMe service for the SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "nvme":{"allowed":"false"}}'
```

- Creates an SVM, allows and configures the NFS service for the SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "nfs":{"allowed":"true",
"enabled":true}}'
```

- Create an SVM and set the max volume limit for the SVM.

```
POST "/api/svm/svms/" '{"name":"testVs", "max_volumes":"200"}'
```

- Creates an SVM and disallows the NDMP service for the SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "ndmp":{"allowed":"false"}}'
```

- Creates an SVM and specifies whether file system analytics is enabled on all newly created volumes in the SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "auto_enable_analytics":true}}'
```

- Creates an SVM and specifies whether volume_activity_tracking is enabled on all newly created volumes in the SVM.

```
POST "/api/svm/svms" '{"name":"testVs",
"auto_enable_activity_tracking":true}}'
```

- Creates an SVM and specifies whether file system analytics is enabled on all newly created volumes in the SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "auto_enable_analytics":true}}'
```

- Creates an SVM and specifies the maximum storage limit for a single SVM.

```
POST "/api/svm/svms" '{"name":"testVs", "storage": {"limit":"4GB"}}'
```

- Creates an SVM and specifies at what percentage of storage capacity an alert message is sent. Default value is 90.

```
POST "/api/svm/svms" '{"name":"testVs", "storage": {"limit":"20GB",
"limit_threshold_alert":"95"}}'
```

- Creates an SVM and specifies the QoS policy group template to be assigned to the SVM.

```
POST "/api/svm/svms" '{"name":"testVs",
"qos_policy_group_template":{"name":"performance-fixed"}}'
```

- Creates an SVM and specifies the QoS adaptive policy group template to be assigned to the SVM.

```
POST "/api/svm/svms" '{"name":"testVs",
"qos_adaptive_policy_group_template":{"name":"performance"}}'
```

**Learn more**

- DOC /svm/svms

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |
| cifs | cifs | |
| comment | string | Comment |
| dns | dns | |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| ipspace | ipspace | Either the UUID or name may be supplied on input. |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| ldap | ldap | |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |
| nis | nis | |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| routes | array[network_route_for_svm] | Optional array of routes for the SVM |
| s3 | s3 | |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| storage | storage | |
| subtype | string | SVM subtype. The SVM subtype sync_destination is created automatically when an SVM of subtype sync_source is created on the source MetroCluster cluster. A POST request with sync_destination as SVM subtype is invalid. SVM of subtype data_engine cannot be explicitly created by the admin and most management changes are not allowed on it. |
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

**Example request**

```
{
  "aggregates": [
    {
      "available_size": 10156560384,
      "name": "aggr1",
      "snaplock_type": "string",
      "state": "string",
      "type": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "anti_ransomware_auto_switch_minimum_incoming_data": "string",
  "anti_ransomware_default_volume_state": "string",
  "anti_ransomware_incoming_write_threshold": "string",
  "anti_ransomware_incoming_write_threshold_percent": "string",
  "cifs": {
    "ad_domain": {
      "default_site": "string",
      "fqdn": "example.com",
      "organizational_unit": "string",
      "password": "string",
      "user": "string"
    },
    "auth-style": "domain",
    "domain_workgroup": "string",
    "name": "CIFS1",
    "workgroup": "workgrp1"
  },
  "comment": "string",
  "dns": {
    "domains": [
      "example.com",
      "example2.example3.com"
    ],
    "servers": [
      "10.224.65.20",
      "2001:db08:a0b:12f0::1"
    ]
  },
  "fc_interfaces": [
    {
      "data_protocol": "string",
      "location": {
        "port": {
```

```
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ip_interfaces": [
    {
      "ip": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "location": {
        "broadcast_domain": {
          "name": "bd1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "home_node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "home_port": {
          "name": "e1b",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "name": "lif1",
      "service_policy": "string",
      "services": [
        "data_nfs"
      ],
      "subnet": {
        "name": "subnet1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
  },
  "language": "c.utf_8",
  "ldap": {
    "ad_domain": "string",
    "base_dn": "string",
    "bind_dn": "string",
    "servers": [
      "string"
    ]
  },
  "max_volumes": "string",
  "name": "svm1",
  "nis": {
    "domain": "string",
    "servers": [
      "string"
    ]
  },
  "nsswitch": {
    "group": [
      "string"
    ],
    "hosts": [
      "string"
    ],
    "namemap": [
      "string"
    ],
    "netgroup": [
      "string"
    ],
    "passwd": [
      "string"
    ]
  },
  "number_of_volumes_in_recovery_queue": 0,
  "qos_adaptive_policy_group_template": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
```

```
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qos_policy_group_template": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "routes": [
    {
      "destination": {
        "address": "10.10.10.7",
```

```
          "netmask": "24"
        },
        "gateway": "10.1.1.1"
      }
    ],
    "s3": {
      "certificate": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "default_unix_user": "string",
      "default_win_user": "string",
      "name": "s3-server-1"
    },
    "snapmirror": {
      "protected_consistency_group_count": 0,
      "protected_volumes_count": 0
    },
    "snapshot_policy": {
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "storage": {
      "allocated": 0,
      "available": 0,
      "used_percentage": 0
    },
    "subtype": "string",
    "total_volume_size_in_recovery_queue": 0,
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
```


**Response**

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 2621580 | Cannot specify options other than SVM name, comment and ipspace for a Vserver that is being configured as the destination for SVM DR. |
| 2621634 | "sync-source" SVM can only be created in a MetroCluster configuration. |
| 2621657 | "sync-destination" SVM can only be created by the system. |
| 13434884 | Cannot create an SVM because of incorrect fields. |
| 13434885 | Non-UTF8 language(s) not supported. |
| 13434888 | IPspace UUID and IPspace name mismatch. |
| 13434889 | Internal Error. Wait and retry. |
| 13434894 | Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again. |
| 13434908 | Invalid SVM name. The name is already in use by another SVM, IPSpace or cluster. |

| Error codes | Description |
| --- | --- |
| 13434909 | Internal Error. Failed to identify the aggregate to host SVM root volume. |
| 13434910 | Internal Error. Failed to allocate new SVM ID. |
| 13434911 | Invalid SVM name. Maximum supported length is 41 if SVM is of type \"sync-source\", otherwise 47. |
| 13434912 | Failed to find IPspace. |
| 13434913 | Internal error: Failed to check if an SVM create operation is in progress. Contact technical support for assistance. |
| 13434914 | Request to create the root volume of the SVM failed because there is not enough space in specified aggregate. |
| 13434915 | Failed to unlock the SVM because SVM create or delete job is in progress. Wait a few minutes, and then try the command again. |
| 13434916 | SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 13434917 | SVM creation successful. |
| 13434918 | IPspace name not provided for creating an SVM. |
| 458753 | Destination and gateway must belong to the same address family. |
| 23724038 | Invalid source for the provided ns-switch database. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

| Name | Type | Description |
|---|---|---|
| available_size | integer | Space available, in bytes. |
| name | string | |
| snaplock_type | string | SnapLock type. |
| state | string | Aggregate state. |
| type | string | Type of aggregate. |
| uuid | string | |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

arw_vserver

Anti-ransomware related information for the SVM.

| Name | Type | Description |
|---|---|---|
| event_log | event_log | |

certificate

Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ad_domain

| Name | Type | Description |
| --- | --- | --- |
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. |
| password | string | The account password used to add this CIFS server to the Active Directory. This is not audited. Valid in POST only. |
| user | string | The user account used to add this CIFS server to the Active Directory. Valid in POST only. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | ad_domain | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the CIFS service. If it is false, only the cluster administrator can manage the service. |
| auth-style | string | Authentication type. |

| Name | Type | Description |
| --- | --- | --- |
| domain_workgroup | string | The NetBIOS name of the domain or workgroup associated with the CIFS server. |
| enabled | boolean | If allowed, setting to true enables the CIFS service. |
| name | string | The NetBIOS name of the CIFS server. |
| workgroup | string | The workgroup name. |

dns

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |

| Name | Type | Description |
| --- | --- | --- |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |

node

The node on which the FC port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the FC port is located. |

fc_port_reference

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the Fibre Channel interface is defined by the location of its port.

| Name | Type | Description |
| --- | --- | --- |
| port | fc_port_reference | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

fc_interface_svm

| Name | Type | Description |
| --- | --- | --- |
| data_protocol | string | The data protocol for which the Fibre Channel interface is configured. |

| Name | Type | Description |
| --- | --- | --- |
| location | location | The location of the Fibre Channel interface is defined by the location of its port. |
| name | string | The name of the Fibre Channel interface. |
| uuid | string | The unique identifier of the Fibre Channel interface. |

fcp

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the FCP service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the FCP service. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

port_svm

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

location

Home_node is optional.

| Name | Type | Description |
|---|---|---|
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. |
| home_node | home_node | |
| home_port | port_svm | |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the subnet. If only the name is provided, the IPspace scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

ip_interface_svm

Interface parameters. Name and home_node are optional.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| location | location | Home_node is optional. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the interface (optional). |
| service_policy | string | Built-in service policies for SVMs. |
| services | array[string] | The services associated with the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Either the UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

iscsi

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the iSCSI service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the iSCSI service. |

ldap

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST. |
| base_dn | string | Specifies the default base DN for all searches. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. SVM API supports anonymous binding. For Simple and SASL LDAP binding, use the LDAP API endpoint. |
| enabled | boolean | Enable LDAP? Setting to true creates a configuration if not already created. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| servers | array[string] | |

ndmp

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NDMP service. If it is false, only the cluster administrator can manage the service. |

nfs

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NFS service. If it is false, only the cluster administrator can manage the service. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | If allowed, setting to true enables the NFS service. |

nis

| Name | Type | Description |
| --- | --- | --- |
| domain | string | The NIS domain to which this configuration belongs. |
| enabled | boolean | Enable NIS? Setting to true creates a configuration if not already created. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

nsswitch

Name service switch configuration

| Name | Type | Description |
| --- | --- | --- |
| group | array[string] | Group sources |
| hosts | array[string] | Host sources |
| namemap | array[string] | NameMap sources |
| netgroup | array[string] | NetGroup sources |
| passwd | array[string] | Password sources |

nvme

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NVMe service. If it is false, only the cluster administrator can manage the service. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | If allowed, setting to true enables the NVMe service. |

qos_adaptive_policy_group_template

This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy

This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy_group_template

This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs).

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

ip_info

IP information

| Name | Type | Description |
|---|---|---|
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

network_route_for_svm

| Name | Type | Description |
|---|---|---|
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
|---|---|---|
| name | string | Certificate name |
| uuid | string | Certificate UUID |

s3

| Name | Type | Description |
|---|---|---|
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |

| Name | Type | Description |
| --- | --- | --- |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether or not to enable S3. Setting this value to true creates a service if one is not yet created. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server. |
| name | string | Specifies the name of the S3 server. A server name length can range from 1 to 253 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. |

snapmirror

Specifies attributes for SVM DR protection.

| Name | Type | Description |
| --- | --- | --- |
| is_protected | boolean | Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data. |
| protected_consistency_group_count | integer | Specifies the number of SVM DR protected consistency groups in the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| protected_volumes_count | integer | Specifies the number of SVM DR protected volumes in the SVM. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

storage

| Name | Type | Description |
| --- | --- | --- |
| allocated | integer | Total size of the volumes in SVM, in bytes. |
| available | integer | Currently available storage capacity in SVM, in bytes. |
| limit | integer | Maximum storage permitted on a single SVM, in bytes. |
| limit_threshold_alert | integer | Indicates at what percentage of storage capacity an alert message is sent. The default value is 90. |
| limit_threshold_exceeded | boolean | Indicates whether the total storage capacity exceeds the alert percentage. |
| used_percentage | integer | The percentage of storage capacity used. |

svm

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |

| Name | Type | Description |
| --- | --- | --- |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |
| cifs | cifs | |
| comment | string | Comment |
| dns | dns | |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| ipspace | ipspace | Either the UUID or name may be supplied on input. |

| Name | Type | Description |
| --- | --- | --- |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| ldap | ldap | |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |
| nis | nis | |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |

| Name | Type | Description |
| --- | --- | --- |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| routes | array[network_route_for_svm] | Optional array of routes for the SVM |
| s3 | s3 | |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| storage | storage | |
| subtype | string | SVM subtype. The SVM subtype sync_destination is created automatically when an SVM of subtype sync_source is created on the source MetroCluster cluster. A POST request with sync_destination as SVM subtype is invalid. SVM of subtype data_engine cannot be explicitly created by the admin and most management changes are not allowed on it. |

| Name | Type | Description |
|---|---|---|
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete an SVM**

DELETE `/svm/svms/{uuid}`

**Introduced In:** 9.6

Deletes an SVM. As a prerequisite, SVM objects must be deleted first. SnapMirror relationships must be deleted and data volumes must be offline and deleted.

- The number of parallel SVMs that can be created must not be greater than five.
- If a sixth SVM POST request is issued, the following error message is generated: "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again."

#### Related ONTAP commands

- `vserver delete`

#### Example

Deleting an individual SVM in the cluster.

```
DELETE "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

#### Learn more

- DOC /svm/svms

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Filter by UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 13434894 | Maximum allowed SVM jobs exceeded. Wait and retry. |
| 2621525 | SVM cannot be deleted as it is associated with an Active Directory configured CIFS server. Delete the CIFS server using "cifs delete" and retry the operation. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve SVM properties**

GET `/svm/svms/{uuid}`

**Introduced In:** 9.6

Retrieves the properties for an individual SVM. This includes protocol configurations such as CIFS and NFS, export policies, name service configurations, and network services.

#### Important notes

- The SVM object includes a large set of fields and can be expensive to retrieve.
- REST APIs only expose a data SVM as an SVM.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `snapmirror.*`

#### Example

Retrieving an individual SVM in the cluster

```
GET "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Filter by UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests. |
| cifs | cifs | |
| comment | string | Comment |
| dns | dns | |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| ipspace | ipspace | Either the UUID or name may be supplied on input. |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| ldap | ldap | |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |
| nis | nis | |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |
| qos_policy | qos_policy | This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| s3 | s3 | |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| state | string | SVM State |
| storage | storage | |
| subtype | string | SVM subtype. The SVM subtype sync_destination is created automatically when an SVM of subtype sync_source is created on the source MetroCluster cluster. A POST request with sync_destination as SVM subtype is invalid. SVM of subtype data_engine cannot be explicitly created by the admin and most management changes are not allowed on it. |
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "available_size": 10156560384,
      "name": "aggr1",
      "snaplock_type": "string",
      "state": "string",
      "type": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "anti_ransomware_auto_switch_minimum_incoming_data": "string",
  "anti_ransomware_default_volume_state": "string",
  "anti_ransomware_incoming_write_threshold": "string",
  "anti_ransomware_incoming_write_threshold_percent": "string",
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cifs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_domain": {
      "default_site": "string",
      "fqdn": "example.com",
      "organizational_unit": "string"
```

```
      },
      "auth-style": "domain",
      "domain_workgroup": "string",
      "name": "CIFS1",
      "workgroup": "workgrp1"
    },
    "comment": "string",
    "dns": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "domains": [
        "example.com",
        "example2.example3.com"
      ],
      "servers": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
      ]
    },
    "fc_interfaces": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "data_protocol": "string",
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "fcp": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "ip_interfaces": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "services": [
        "data_nfs"
      ],
      "subnet": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "subnet1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "iscsi": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "language": "c.utf_8",
  "ldap": {
    "ad_domain": "string",
    "base_dn": "string",
    "bind_dn": "string",
    "servers": [
      "string"
    ]
  },
```

```
  "max_volumes": "string",
  "name": "svm1",
  "nfs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "nis": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "domain": "string",
    "servers": [
      "string"
    ]
  },
  "nsswitch": {
    "group": [
      "string"
    ],
    "hosts": [
      "string"
    ],
    "namemap": [
      "string"
    ],
    "netgroup": [
      "string"
    ],
    "passwd": [
      "string"
    ]
  },
  "number_of_volumes_in_recovery_queue": 0,
  "nvme": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "qos_adaptive_policy_group_template": {
```

```
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
```

```
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qos_policy_group_template": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
```

```
  "s3": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "default_unix_user": "string",
    "default_win_user": "string",
    "name": "s3-server-1"
  },
  "snapmirror": {
    "protected_consistency_group_count": 0,
    "protected_volumes_count": 0
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "storage": {
    "allocated": 0,
    "available": 0,
    "used_percentage": 0
  },
  "subtype": "string",
  "total_volume_size_in_recovery_queue": 0,
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
|---|---|
| 262188 | Field "top_metric" was specified twice (to "iops.read" and "throughput.read"). |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

aggregates

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| available_size | integer | Space available, in bytes. |
| name | string | |
| snaplock_type | string | SnapLock type. |
| state | string | Aggregate state. |
| type | string | Type of aggregate. |
| uuid | string | |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

arw_vserver

Anti-ransomware related information for the SVM.

| Name | Type | Description |
|---|---|---|
| event_log | event_log | |

certificate

Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ad_domain

| Name | Type | Description |
|---|---|---|
| default_site | string | The default site used by LIFs that do not have a site membership. |
| fqdn | string | The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. |
| organizational_unit | string | Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. |

cifs

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ad_domain | ad_domain | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the CIFS service. If it is false, only the cluster administrator can manage the service. |
| auth-style | string | Authentication type. |
| domain_workgroup | string | The NetBIOS name of the domain or workgroup associated with the CIFS server. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | If allowed, setting to true enables the CIFS service. |
| name | string | The NetBIOS name of the CIFS server. |
| workgroup | string | The workgroup name. |

dns

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |

node

The node on which the FC port is located.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the node on which the FC port is located. |

fc_port_reference

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the Fibre Channel interface is defined by the location of its port.

| Name | Type | Description |
|---|---|---|
| port | fc_port_reference | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

fc_interface_svm

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| data_protocol | string | The data protocol for which the Fibre Channel interface is configured. |
| name | string | The name of the Fibre Channel interface. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the Fibre Channel interface. |

fcp

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the FCP service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the FCP service. |

ip

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

port_svm

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

location

Home_node is optional.

| Name | Type | Description |
|---|---|---|
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. |
| home_node | home_node | |
| home_port | port_svm | |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

ip_interface_svm

Interface parameters. Name and home_node are optional.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ip | ip | IP information |
| name | string | The name of the interface (optional). |

| Name | Type | Description |
| --- | --- | --- |
| services | array[string] | The services associated with the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Either the UUID or name may be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

iscsi

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the iSCSI service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the iSCSI service. |

ldap

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST. |
| base_dn | string | Specifies the default base DN for all searches. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. SVM API supports anonymous binding. For Simple and SASL LDAP binding, use the LDAP API endpoint. |
| enabled | boolean | Enable LDAP? Setting to true creates a configuration if not already created. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| servers | array[string] | |

ndmp

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NDMP service. If it is false, only the cluster administrator can manage the service. |

nfs

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NFS service. If it is false, only the cluster administrator can manage the service. |

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | If allowed, setting to true enables the NFS service. |

nis

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| domain | string | The NIS domain to which this configuration belongs. |
| enabled | boolean | Enable NIS? Setting to true creates a configuration if not already created. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

nsswitch

Name service switch configuration

| Name | Type | Description |
| --- | --- | --- |
| group | array[string] | Group sources |
| hosts | array[string] | Host sources |
| namemap | array[string] | NameMap sources |
| netgroup | array[string] | NetGroup sources |
| passwd | array[string] | Password sources |

nvme

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NVMe service. If it is false, only the cluster administrator can manage the service. |
| enabled | boolean | If allowed, setting to true enables the NVMe service. |

qos_adaptive_policy_group_template

This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy

This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy_group_template

This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs).

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| family | string | IPv4 or IPv6 |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

network_route_for_svm

| Name | Type | Description |
| --- | --- | --- |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

s3

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. |
| certificate | certificate | Specifies the certificate that will be used for creating HTTPS connections to the S3 server. |
| default_unix_user | string | Specifies the default UNIX user for NAS Access. |
| default_win_user | string | Specifies the default Windows user for NAS Access. |
| enabled | boolean | Specifies whether or not to enable S3. Setting this value to true creates a service if one is not yet created. |
| is_http_enabled | boolean | Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server. |
| is_https_enabled | boolean | Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server. |
| name | string | Specifies the name of the S3 server. A server name length can range from 1 to 253 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-". |
| port | integer | Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80. |
| secure_port | integer | Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443. |

snapmirror

Specifies attributes for SVM DR protection.

| Name | Type | Description |
|---|---|---|
| is_protected | boolean | Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data. |
| protected_consistency_group_count | integer | Specifies the number of SVM DR protected consistency groups in the SVM. |
| protected_volumes_count | integer | Specifies the number of SVM DR protected volumes in the SVM. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

storage

| Name | Type | Description |
|---|---|---|
| allocated | integer | Total size of the volumes in SVM, in bytes. |
| available | integer | Currently available storage capacity in SVM, in bytes. |
| limit | integer | Maximum storage permitted on a single SVM, in bytes. |
| limit_threshold_alert | integer | Indicates at what percentage of storage capacity an alert message is sent. The default value is 90. |
| limit_threshold_exceeded | boolean | Indicates whether the total storage capacity exceeds the alert percentage. |

| Name | Type | Description |
|---|---|---|
| used_percentage | integer | The percentage of storage capacity used. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update SVM properties

PATCH `/svm/svms/{uuid}`

**Introduced In:** 9.6

Updates one or more of the following properties of an individual SVM: SVM name, SVM default volume language code, SVM comment, and SVM state.

#### Related ONTAP commands

- `vserver modify`
- `vserver rename`
- `vserver start`
- `vserver stop`
- `security ssl modify`
- `vserver add-protocols`
- `vserver remove-protocols`

**Examples**

1. Stops an SVM and updates the "comment" field for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"state":"stopped", "comment":"This SVM is stopped."}'
```

1. Starts an SVM and updates the "comment" field for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"state":"running", "comment":"This SVM is running."}'
```

1. Updates the "language" field for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"language":"en.UTF-8"}'
```

1. Updates the "name" field for an SVM or renames the SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"name":"svm_new"}'
```

1. Updates the aggregates for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"aggregates":["{"name":aggr1"},{"name":"aggr2"},{"name":"aggr3"}]}'
```

1. Updates the snapshot policy for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"snapshot_policy":{"name":"custom1"}}'
```

1. Updates the TLS certificate for an individual SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"certificate":{"uuid":"1cd8a442-86d1-11e0-ae1c-123478563412"}}'
```

1. Updates the QoS policy for the SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"qos_policy_group":{"name":"qpolicy1"}}'
```

1. Allows NFS protocol which was previously disallowed for the SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"nfs":{"allowed":"true"}}'
```

1. Updates the max volume limit for the SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"max_volumes":"200"}'
```

1. Updates whether file system analytics is enabled on all newly created volumes in the SVM.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"auto_enable_analytics":"true"}'
```

1. Updates whether volume activity tracking is enabled on all newly created volumes in the SVM.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"auto_enable_activity_tracking":"true"}'
```

1. Updates the QoS adaptive policy group template for the SVM.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"qos_adaptive_policy_group_template":{"name":"aqpolicy1"}}'
```

1. Updates the maximum storage permitted on a single SVM.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"storage":{"limit":"40GB"}}'
```

1. Updates the percentage of storage capacity at which an alert message is sent.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"storage":{"limit":"400MB", "limit_threshold_alert":"98"}}'
```

1. Updates the QoS policy group template for the SVM.

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"qos_policy_group_template":{"name":"policy1"}}'
```

1. Updates the S3 protocol that was previously disallowed for the SVM

```
  PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"s3":{"allowed":"true"}}'
```

**Learn more**

- DOC /svm/svms

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Filter by UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests. |
| cifs | cifs | |
| comment | string | Comment |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |

| Name | Type | Description |
| --- | --- | --- |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |
| qos_policy | qos_policy | This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| s3 | s3 | |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| state | string | SVM State |
| storage | storage | |

| Name | Type | Description |
| --- | --- | --- |
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

**Example request**


```
{
  "aggregates": [
    {
      "available_size": 10156560384,
      "name": "aggr1",
      "snaplock_type": "string",
      "state": "string",
      "type": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "anti_ransomware_auto_switch_minimum_incoming_data": "string",
  "anti_ransomware_default_volume_state": "string",
  "anti_ransomware_incoming_write_threshold": "string",
  "anti_ransomware_incoming_write_threshold_percent": "string",
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cifs": {
    "auth-style": "domain",
    "domain_workgroup": "string",
    "enabled": true,
    "workgroup": "workgrp1"
  },
  "comment": "string",
  "fc_interfaces": [
    {
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "fcp": {
    "enabled": true
  },
  "ip_interfaces": [
    {
      "services": [
        "data_nfs"
      ],
      "subnet": {
        "name": "subnet1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      }
    ],
    "iscsi": {
      "enabled": true
    },
    "language": "c.utf_8",
    "max_volumes": "string",
    "name": "svm1",
    "nfs": {
      "enabled": true
    },
    "nsswitch": {
      "group": [
        "string"
      ],
      "hosts": [
        "string"
      ],
      "namemap": [
        "string"
      ],
      "netgroup": [
        "string"
      ],
      "passwd": [
        "string"
      ]
    },
    "number_of_volumes_in_recovery_queue": 0,
    "nvme": {
      "enabled": true
    },
    "qos_adaptive_policy_group_template": {
      "max_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput": [
        "900KB/s",
```

```
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qos_policy_group_template": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
```

```
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "s3": {
    "enabled": true,
    "is_http_enabled": true,
    "is_https_enabled": true
  },
  "snapmirror": {
    "protected_consistency_group_count": 0,
    "protected_volumes_count": 0
  },
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "storage": {
    "allocated": 0,
    "available": 0,
    "used_percentage": 0
  },
  "total_volume_size_in_recovery_queue": 0,
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```


**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error codes | Description |
| --- | --- |
| 13434880 | Failed to modify SVM parameters. |
| 13434881 | Failed to rename SVM. |
| 13434883 | SVM parameters except name modified successfully. |
| 13434885 | Non-UTF8 language(s) not supported. |
| 13434886 | Invalid snapshot policy. |
| 13434902 | Modification of NSSwitch parameters failed for the SVM. |
| 13434906 | Operation not supported for an SVM of type sync-destination. |
| 12451843 | Certificate does not exist. |
| 13434908 | Invalid SVM name. The name is already in use by another SVM, IPSpace or cluster. |
| 13434916 | SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 13434915 | Failed to unlock the SVM because SVM create or delete job is in progress. Wait a few minutes, and then try the command again. |

| Error codes | Description |
| --- | --- |
| 13434911 | Invalid SVM name. Maximum supported length is 41 if SVM is of type \"sync-source\", otherwise 47. |
| 262179 | Unexpected argument "storage_limit". |
| 23724038 | Invalid source for the provided ns-switch database. |
| 2621779 | Operation not supported for an SVM of type data-engine. |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

aggregates

| Name | Type | Description |
|---|---|---|
| available_size | integer | Space available, in bytes. |
| name | string | |
| snaplock_type | string | SnapLock type. |
| state | string | Aggregate state. |
| type | string | Type of aggregate. |
| uuid | string | |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

arw_vserver

Anti-ransomware related information for the SVM.

| Name | Type | Description |
|---|---|---|
| event_log | event_log | |

certificate

Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

ad_domain

| Name | Type | Description |
| --- | --- | --- |
| default_site | string | The default site used by LIFs that do not have a site membership. |

cifs

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the CIFS service. If it is false, only the cluster administrator can manage the service. |
| auth-style | string | Authentication type. |
| domain_workgroup | string | The NetBIOS name of the domain or workgroup associated with the CIFS server. |
| workgroup | string | The workgroup name. |

dns

| Name | Type | Description |
| --- | --- | --- |
| domains | array[string] | A list of DNS domains.<br>Domain names have the following requirements:<br>• The name must contain only the following characters: A through Z,<br>a through z, 0 through 9, ".", "-" or "_".<br>• The first character of each label, delimited by ".", must be one<br>of the following characters: A through Z or a through z or 0 through 9.<br>• The last character of each label, delimited by ".", must be one of<br>the following characters: A through Z, a through z, or 0 through 9.<br>• The top level domain must contain only the following characters: A<br>through Z, a through z.<br>• The system reserves the following names:"all", "local", and "localhost". |
| servers | array[string] | The list of IP addresses of the DNS servers. Addresses can be either<br>IPv4 or IPv6 addresses. |

node

The node on which the FC port is located.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the node on which the FC port is located. |

fc_port_reference

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the FC port. |
| node | node | The node on which the FC port is located. |
| uuid | string | The unique identifier of the FC port. |

location

The location of the Fibre Channel interface is defined by the location of its port.

| Name | Type | Description |
| --- | --- | --- |
| port | fc_port_reference | An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network. |

fc_interface_svm

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the Fibre Channel interface. |

fcp

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the FCP service. If it is false, only the cluster administrator can manage the service. |

ip

IP information

broadcast_domain

Broadcast domain UUID along with a readable name.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the broadcast domain, scoped to its IPspace |
| uuid | string | Broadcast domain UUID |

home_node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

port_svm

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

location

Home_node is optional.

| Name | Type | Description |
| --- | --- | --- |
| broadcast_domain | broadcast_domain | Broadcast domain UUID along with a readable name. |
| home_node | home_node | |
| home_port | port_svm | |

ip_subnet_reference

A named subnet. Either UUID or name can be supplied on input.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the subnet. If only the name is provided, the IPspace scope<br>must be provided by the object this object is embedded in. |
| uuid | string | The UUID that uniquely identifies the subnet. |

ip_interface_svm

Interface parameters. Name and home_node are optional.

| Name | Type | Description |
|---|---|---|
| ip | ip | IP information |
| services | array[string] | The services associated with the interface. |
| subnet | ip_subnet_reference | A named subnet. Either UUID or name can be supplied on input. |
| uuid | string | The UUID that uniquely identifies the interface. |

ipspace

Either the UUID or name may be supplied on input.

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

iscsi

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
|---|---|---|
| allowed | boolean | If this is set to true, an SVM administrator can manage the iSCSI service. If it is false, only the cluster administrator can manage the service. |

ldap

| Name | Type | Description |
|---|---|---|
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST. |

| Name | Type | Description |
| --- | --- | --- |
| base_dn | string | Specifies the default base DN for all searches. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. SVM API supports anonymous binding. For Simple and SASL LDAP binding, use the LDAP API endpoint. |
| enabled | boolean | Enable LDAP? Setting to true creates a configuration if not already created. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| servers | array[string] | |

ndmp

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NDMP service. If it is false, only the cluster administrator can manage the service. |

nfs

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the NFS service. If it is false, only the cluster administrator can manage the service. |

nis

| Name | Type | Description |
| --- | --- | --- |
| domain | string | The NIS domain to which this configuration belongs. |

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Enable NIS? Setting to true creates a configuration if not already created. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

nsswitch

Name service switch configuration

| Name | Type | Description |
|---|---|---|
| group | array[string] | Group sources |
| hosts | array[string] | Host sources |
| namemap | array[string] | NameMap sources |
| netgroup | array[string] | NetGroup sources |
| passwd | array[string] | Password sources |

nvme

Available for GET, POST, and PATCH requests.

| Name | Type | Description |
|---|---|---|
| allowed | boolean | If this is set to true, an SVM administrator can manage the NVMe service. If it is false, only the cluster administrator can manage the service. |

qos_adaptive_policy_group_template

This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy

This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos_policy_group_template

This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs).

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

ip_info

IP information

| Name | Type | Description |
| --- | --- | --- |
| address | string | IPv4 or IPv6 address |
| netmask | string | Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the default value is 64 with a valid range of 1 to 127. Output is always the netmask length. |

network_route_for_svm

| Name | Type | Description |
| --- | --- | --- |
| destination | ip_info | IP information |
| gateway | string | The IP address of the gateway router leading to the destination. |

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

| Name | Type | Description |
| --- | --- | --- |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

s3

| Name | Type | Description |
| --- | --- | --- |
| allowed | boolean | If this is set to true, an SVM administrator can manage the S3 service. If it is false, only the cluster administrator can manage the service. |

snapmirror

Specifies attributes for SVM DR protection.

| Name | Type | Description |
| --- | --- | --- |
| is_protected | boolean | Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data. |
| protected_consistency_group_count | integer | Specifies the number of SVM DR protected consistency groups in the SVM. |
| protected_volumes_count | integer | Specifies the number of SVM DR protected volumes in the SVM. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

storage

| Name | Type | Description |
| --- | --- | --- |
| allocated | integer | Total size of the volumes in SVM, in bytes. |
| available | integer | Currently available storage capacity in SVM, in bytes. |
| limit | integer | Maximum storage permitted on a single SVM, in bytes. |
| limit_threshold_alert | integer | Indicates at what percentage of storage capacity an alert message is sent. The default value is 90. |
| limit_threshold_exceeded | boolean | Indicates whether the total storage capacity exceeds the alert percentage. |
| used_percentage | integer | The percentage of storage capacity used. |

svm

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates. |
| aggregates_delegated | boolean | This property is true when the administrator has delegated the aggregates for the SVM volumes. |
| anti_ransomware | arw_vserver | Anti-ransomware related information for the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_duration_without_new_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_from_learning_to_enabled | boolean | This property specifies whether anti-ransomware state of the volumes in this SVM are automatically switched by the system from “learning” (dry-run) to “enabled” (active) state after sufficient learning. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_file_count | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have a minimum file count in “learning” state. This parameter optionally specifies the minimum number of newly created files in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_file_extension | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have minimum number of file extensions in “learning” state. This parameter optionally specifies the minimum number of new file extensions in “learning” state in a given volume to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_auto_switch_minimum_incoming_data | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_auto_switch_minimum_learning_period | integer | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. This field will no longer be supported in a future release. |
| anti_ransomware_default_volume_state | string | Specifies the default anti-ransomware state of the volumes in the SVM. The default "anti_ransomware_default_volume_state" property is disabled for POST operations. If this value is "disabled", anti-ransomware protection is disabled by default on the new volumes that are created in the SVM. If this value is "dry_run", anti-ransomware protection is in learning mode by default on the new volumes that are created in the SVM. When the anti-ransomware license is not present, this property is ignored and volumes will be created with the "disabled" state. This value "dry_run" will no longer be supported in a future release. |

| Name | Type | Description |
| --- | --- | --- |
| anti_ransomware_incoming_write_threshold | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in GB) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field is no longer supported. |
| anti_ransomware_incoming_write_threshold_percent | string | One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this SVM from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. This field will no longer be supported in a future release. |
| auto_enable_activity_tracking | boolean | Specifies whether volume activity tracking is automatically enabled on volumes that are created in the SVM. |
| auto_enable_analytics | boolean | Specifies whether file system analytics is automatically enabled on volumes that are created in the SVM. |

| Name | Type | Description |
| --- | --- | --- |
| certificate | certificate | Support for this field will be removed in a future release. Please use /svm/svms/{svm.uuid}/web for this field. Certificate for incoming TLS connection requests. |
| cifs | cifs | |
| comment | string | Comment |
| fc_interfaces | array[fc_interface_svm] | FC Interface for the SVM |
| fcp | fcp | Available for GET, POST, and PATCH requests. |
| ip_interfaces | array[ip_interface_svm] | IP interfaces for the SVM |
| is_space_enforcement_logical | boolean | Indicates whether logical space enforcement for the SVM is enabled. |
| is_space_reporting_logical | boolean | Indicates whether logical space reporting for the SVM is enabled. |
| iscsi | iscsi | Available for GET, POST, and PATCH requests. |
| language | string | Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests. |
| max_volumes | string | This property is used by cluster administrator to specify the limit on maximum number of volumes allowed in the SVM. The value can be either the string "unlimited" or a number. |
| name | string | The name of the SVM. |
| ndmp | ndmp | |
| nfs | nfs | |

| Name | Type | Description |
| --- | --- | --- |
| nsswitch | nsswitch | Name service switch configuration |
| number_of_volumes_in_recovery_queue | integer | Number of volumes in the recovery queue. |
| nvme | nvme | Available for GET, POST, and PATCH requests. |
| qos_adaptive_policy_group_template | qos_adaptive_policy_group_template | This optionally specifies which QoS adaptive policy group to apply to the SVM as a template. This policy group will then be assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. |
| qos_policy | qos_policy | This optionally specifies which QoS policy group to apply to the Vserver. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. |
| qos_policy_group_template | qos_policy_group_template | This optionally specifies which QoS non-shared policy group to apply to the SVM as a template. This policy group is then assigned to volumes created or rehosted into this SVM. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs). |
| s3 | s3 | |
| snapmirror | snapmirror | Specifies attributes for SVM DR protection. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | SVM State |
| storage | storage | |
| total_volume_size_in_recovery_queue | integer | Sum of the sizes of the volumes in the recovery queue. |
| uuid | string | The unique identifier of the SVM. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve clients

### Retrieve clients

#### Overview

You can use this API to retrieve a list of clients with the most I/O activity for FlexVol and FlexGroup volumes belonging to a specified SVM, within the past several seconds. To obtain this list, only the volumes which have

the activity tracking feature enabled are considered.

This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` and `throughput` metrics. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API supports returning only one I/O activity type per request.

#### Approximate accounting and error bars

When too many clients have had recent activity, some clients may be dropped from the list. In that situation, the spread of values in the `error` field will increase indicating we have larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the clients that would have otherwise been included, may not be present in the final list returned by the API.

#### Enabling and disabling activity tracking feature

The following APIs can be used to enable, disable, and retrieve the activity tracking state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"on"}'

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"off"}'

– GET /api/storage/volumes/{uuid}/?fields=activity_tracking

#### Excluded volumes list

Optionally, the API returns an excluded list of activity tracking-enabled volumes, which were not accounted for when obtaining the list of clients with the most I/O activity for the SVM. This excluded list contains both the volume information and the reason for exclusion.

#### Failure to return list of clients with most I/O activity

The API can sometimes fail to return the list of clients with the most I/O activity, due to the following reasons:

– The volumes belonging to the SVM do not have the activity tracking feature enabled.

– The volumes belonging to the SVM have not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of clients with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Retrieve a list of the clients with the most I/O activity

For a report on the clients with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` I/O activity type into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of clients with the greatest number of average read operations per second. The current maximum number of clients returned by the API for an I/O activity type is 25.

– GET /api/svm/svms/{svm.uuid}/top-metrics/clients

#### Examples

**Retrieving a list of the clients with the greatest average number of write operations per second:**

```
# The API:
GET /api/svm/svms/{svm.uuid}/top-metrics/clients

# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/clients?top_metric=iops.write"

# The Response:
{
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "client_ip": "172.28.71.128"
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "client_ip": "172.28.71.179"
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
```

```
      "write": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "client_ip": "172.28.51.62"
  }
],
"num_records": 3,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol10",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "resource limit exceeded",
      "code": "12345"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc227-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol22",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
```

```
      "code": "23456"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/clients?top_metric=iops.write"
  }
}
}
```

**Example showing the behavior of the API when there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/clients?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for SVM \"vs1\" returned zero
records. Check whether the activity tracking enabled volumes belonging to
the SVM have read/write traffic. Refer to the REST API documentation for
more information on why there might be no records.",
  "code": "124519405",
},
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/clients?top_metric=throughput.write"
  }
}
}
```

### List clients with the most I/O activity

GET `/svm/svms/{svm.uuid}/top-metrics/clients`

**Introduced In:** 9.11

Retrieves a list of clients with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| svm.name | string | query | False | Filter by svm.name |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| throughput.write | integer | query | False | Filter by throughput.write |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.read | integer | query | False | Filter by throughput.read |
| client_ip | string | query | False | Filter by client_ip |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| excluded_volumes | array[top_metrics_svm_client_excluded_volume] | List of volumes that are not included in the SVM activity tracking REST API. |
| notice | notice | Optional field that indicates why no records are returned by the SVM activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_svm_client] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "excluded_volumes": [
    {
      "reason": {
        "code": "111411207",
        "message": "The volume is offline."
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "notice": {
    "code": "111411207",
    "message": "The volume is offline."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "client_ip": "192.168.185.170",
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
```

```
        },
        "read": 5,
        "write": 10
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 12,
        "write": 2
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124519405 | The activity tracking report for SVM svm.name returned zero records. Check whether the activity tracking enabled volumes belonging to the SVM have read/write traffic. Refer to the REST API documentation for more information on why there might be no records. |
| 124519406 | Failed to get the activity tracking report for SVM svm.name. Reason:<Reason for="" failure="">.</Reason> |
| 124519407 | SVM wildcard queries are not supported for activity tracking reports. |

| Error Code | Description |
| --- | --- |
| 124519408 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination for SVM DR. |
| 124519409 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination of a MetroCluster SVM relationship and the SVM admin state is stopped. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

reason

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why the volume is not included in the SVM activity tracking REST API. |
| message | string | Details why the volume is not included in the SVM activity tracking REST API. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_svm_client_excluded_volume

List of volumes and their details as to why they are not included in the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| reason | reason | |
| volume | volume | |

notice

Optional field that indicates why no records are returned by the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

top_metric_value_error_bounds

| Name | Type | Description |
| --- | --- | --- |
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

top_metrics_svm_client

Aggregated information about a client's IO activity at a SVM scope.

| Name | Type | Description |
| --- | --- | --- |
| client_ip | string | IP address of the client. Both IPv4 and IPv6 IP addresses are supported. |
| iops | iops | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve directories

### Retrieve directories

#### Overview

You can use this API to retrieve a list of directories with the most I/O activity for FlexVol and FlexGroup volumes belonging to a specified SVM, within the past several seconds. To obtain this list, only the volumes which have the activity tracking feature enabled are considered.

This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` and `throughput` metrics. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API supports returning only one I/O activity type per request.

#### Enabling and disabling activity tracking feature

The following APIs can be used to enable, disable, and retrieve the activity tracking state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"on"}'

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"off"}'

– GET /api/storage/volumes/{uuid}/?fields=activity_tracking

#### Excluded volumes list

Optionally, the API returns an excluded list of activity tracking-enabled volumes, which were not accounted for when obtaining the list of clients with the most I/O activity for the SVM. This excluded list contains both the volume information and the reason for exclusion.

#### Approximate accounting and error bars

When too many directories have recent activity, some directories might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the directories that would have otherwise been included might not be present in the final list returned by the API.

#### Failure to return list of directories with most I/O activity

The API can sometimes fail to return the list of directories with the most I/O activity, due to the following reasons:

– The volumes belonging to the SVM do not have the activity tracking feature enabled.

– The volumes belonging to the SVM have not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of directories with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return pathnames

The API can sometimes fail to obtain filesystem pathnames for certain directories, either due to internal transient errors or if those directories have been recently deleted.<br> In such cases, instead of the pathname, the API will return "{<volume_instance_uuid>.<fileid>}" for that directory.</fileid></volume_instance_uuid>

#### Retrieve a list of the directories with the most I/O activity

For a report on the directories with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of directories with the greatest number of average read operations per second. The current maximum number of directories returned by the API for an I/O activity type is 25.

– GET /api/svm/svms/{svm.uuid}/top-metrics/directories

#### Examples

**Retrieving a list of the directories with the greatest average number of read operations per second:**

```
# The API:
GET /api/svm/svms/{svm.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/directories?top_metric=iops.read"

# The Response:
{
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "path": "/vol/fv1/dir1/dir2",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef86",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef86"
          }
        }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
```

```
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "path": "/vol/fv2/dir3/dir4",
    "junction-path": "/fv2",
    "volume": {
        "name": "fv2",
        "uuid": "11b293df-e9d7-46cc-a9ce-2df8e52ef811",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/11b293df-e9d7-46cc-a9ce-
2df8e52ef811"
          }
        }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/11b293df-e9d7-46cc-a9ce-
2df8e52ef811/files/dir3%2Fdir4"
      },
      "metadata": {
        "href": "/api/storage/volumes/11b293df-e9d7-46cc-a9ce-
2df8e52ef811/files/dir3%2Fdir4?return_metadata=true"
      }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 345,
      "error": {
```

```
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "path": "/vol/fv1/dir12",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
          }
        }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol10",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 12345
    },
    "_links": {
```

```
      "self": {
        "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc227-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol22",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 23456
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/directories?top_metric=iops.read"
  }
}
}
```


#### Retrieving a list of the directories with the most read traffic, with failure to obtain pathnames

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/directories?top_metric=iops.read"

# The Response:
```

```
{
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "path": "73b293df-e9d7-46cc-a9ce-2df8e52ef86.1232",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef86",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef86"
          }
        }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "path": "11b293df-e9d7-46cc-a9ce-2df8e52ef811.6574",
    "junction-path": "/fv2",
    "volume": {
        "name": "fv2",
        "uuid": "11b293df-e9d7-46cc-a9ce-2df8e52ef811",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/11b293df-e9d7-46cc-a9ce-
2df8e52ef811"
```

```
          }
        }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "path": "73b293df-e9d7-46cc-a9ce-2df8e52ef864.7844",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
          }
        }
    }
  }
],
"num_records": 3,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol10",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 12345
    },
```

```
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc227-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol22",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 23456
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/directories?top_metric=iops.read"
  }
}
}
```


**Example showing the behavior of the API where there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/directories?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for SVM \"vs1\" returned zero
records. Check whether the activity tracking enabled volumes belonging to
the SVM have read/write traffic. Refer to the REST API documentation for
more information on why there might be no records.",
  "code": "124519405",
},
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/directories?top_metric=throughput.write"
  }
}
}
```

---

### List directories with the most I/O activity

GET `/svm/svms/{svm.uuid}/top-metrics/directories`

**Introduced In:** 9.11

Retrieves a list of directories with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| max_records_per_svm | integer | query | False | Max records per svm. |
| path | string | query | False | Filter by path |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |
| svm.name | string | query | False | Filter by svm.name |
| junction-path | string | query | False | Filter by junction-path |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| excluded_volumes | array[top_metrics_svm_directory_excluded_volume] | List of volumes that are not included in the SVM activity tracking REST API. |
| notice | notice | Optional field that indicates why no records are returned by the SVM activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_svm_directory] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "excluded_volumes": [
    {
      "reason": {
        "code": "111411",
        "message": "The volume is offline."
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on svm."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 10,
        "write": 5
      },
      "junction-path": "/fv",
      "path": "/vol/fv/dir_abc/dir_123/dir_20",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 3,
        "write": 20
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124519405 | The activity tracking report for SVM svm.name returned zero records. Check whether the activity tracking enabled volumes belonging to the SVM have read/write traffic. Refer to the REST API documentation for more information on why there might be no records. |
| 124519406 | Failed to get the activity tracking report for SVM svm.name. Reason:<Reason for="" failure="">.</Reason> |
| 124519407 | SVM wildcard queries are not supported for activity tracking reports. |
| 124519408 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination for SVM DR. |
| 124519409 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination of a MetroCluster SVM relationship and the SVM admin state is stopped. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

reason

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why the volume is not included in the SVM activity tracking REST API. |
| message | string | Details why the volume is not included in the SVM activity tracking REST API. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_svm_directory_excluded_volume

List of volumes and their details as to why they are not included in the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| reason | reason | |
| volume | volume | |

notice

Optional field that indicates why no records are returned by the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

_links

| Name | Type | Description |
| --- | --- | --- |
| metadata | href | |
| self | href | |

top_metric_value_error_bounds

| Name | Type | Description |
| --- | --- | --- |
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |

| Name | Type | Description |
|---|---|---|
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

top_metrics_svm_directory

Information about a directory's IO activity.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| iops | iops | |
| junction-path | string | Junction path of the file. |
| path | string | Path of the directory. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Retrieve files

## Retrieve files

### Overview

You can use this API to retrieve a list of files with the most I/O activity for FlexVol and FlexGroup volumes belonging to a specified SVM, within the past several seconds. To obtain this list, only the volumes which have the activity tracking feature enabled are considered.

This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` and `throughput` metrics. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API supports returning only one I/O activity type per request.

### Enabling and disabling activity tracking feature

The following APIs can be used to enable, disable, and retrieve the activity tracking state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"on"}'

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"off"}'

– GET /api/storage/volumes/{uuid}/?fields=activity_tracking

### Excluded volumes list

Optionally, the API returns an excluded list of activity tracking-enabled volumes, which were not accounted for when obtaining the list of clients with the most I/O activity for the SVM. This excluded list contains both the volume information and the reason for exclusion.

### Approximate accounting and error bars

When too many files have recent activity, some files might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the files that would have otherwise been included might not be present in the final list returned by the API.

### Failure to return list of files with most I/O activity

The API can sometimes fail to return the list of files with the most I/O activity, due to the following reasons.

– The volumes belonging to the SVM do not have the activity tracking feature enabled.

– The volumes belonging to the SVM have not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of files with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return pathnames

The API can sometimes fail to obtain the filesystem pathnames for certain files, either due to internal transient errors or if those files have been recently deleted.<br> In such cases, instead of the pathname, the API will return "{<volume_instance_uuid>.<fileid>}" for that file.</fileid></volume_instance_uuid>

#### Retrieve a list of the files with the most I/O activity

For a report on the files with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of the files with the greatest number of the average read operations per second. The current maximum number of files returned by the API for an I/O activity type is 25.

– GET /api/svm/svms/{svm.uuid}/top-metrics/files

#### Examples

**Retrieving a list of the files with the greatest average number of write bytes received per second:**

```
# The API:
GET /api/svm/svms/{svm.uuid}/top-metrics/files

# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/files?top_metric=throughput.write"

# The Response:
{
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "throughput": {
      "write": 24,
      "error": {
        "lower_bound": 24,
        "upper_bound": 29
      }
    },
    "path": "/vol/fv1/d5/f5",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
        "_links": {
          "self": {
```

```
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
          }
        }
    },
    "_links": {
      "metadata": {
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/d5%2Ff5?return_metadata=true"
      }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "throughput": {
      "write": 12,
      "error": {
        "lower_bound": 12,
        "upper_bound": 22
      }
    },
    "path": "/vol/fv2/d6/f6",
    "junction-path": "/fv2",
    "volume": {
        "name": "fv2",
        "uuid": "2ea74c3e-d5ca-11eb-8fbb-005056ac0f33",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/2ea74c3e-d5ca-11eb-8fbb-
005056ac0f33"
          }
        }
    },
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/2ea74c3e-d5ca-11eb-8fbb-
005056ac0f33/files/d6%2Ff6?return_metadata=true"
      }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
```

```
    "throughput": {
      "write": 8,
      "error": {
        "lower_bound": 8,
        "upper_bound": 10
      }
    },
    "path": "/vol/fv3/d3/f3",
    "junction-path": "/fv3",
    "volume": {
        "name": "fv3",
        "uuid": "1ca74c3e-d5ca-11eb-8fbb-005056ac0f88",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/1ca74c3e-d5ca-11eb-8fbb-
005056ac0f88"
          }
        }
    },
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/1ca74c3e-d5ca-11eb-8fbb-
005056ac0f88/files/d3%2Ff3?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc224-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 12345
    },
    "_links": {
      "self": {
```

```
        "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc224-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol2",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 23456
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/top-
metrics/files?top_metric=throughput.write"
  }
}
}
```


---

#### Retrieving a list of the files with the most read traffic, with failure to obtain pathnames

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/files?top_metric=iops.read"

# The Response:
{
```

```
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "path": "73b293df-e9d7-46cc-a9ce-2df8e52ef86.1232",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef86",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef86"
          }
        }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "path": "11b293df-e9d7-46cc-a9ce-2df8e52ef811.6574",
    "junction-path": "/fv2",
    "volume": {
        "name": "fv2",
        "uuid": "11b293df-e9d7-46cc-a9ce-2df8e52ef811",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/11b293df-e9d7-46cc-a9ce-
2df8e52ef811"
          }
```

```
        }
    }
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "read": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "path": "73b293df-e9d7-46cc-a9ce-2df8e52ef864.7844",
    "junction-path": "/fv1",
    "volume": {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
        "_links": {
          "self": {
              "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
          }
        }
    }
  }
],
"num_records": 3,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol10",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 12345
    },
    "_links": {
```

```
      "self": {
        "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc227-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol22",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 23456
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43/top-
metrics/files?top_metric=iops.read"
  }
}
}
```


**Example showing the behavior of the API where there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/files?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for SVM \"vs1\" returned zero
records. Check whether the activity tracking enabled volumes belonging to
the SVM have read/write traffic. Refer to the REST API documentation for
more information on why there might be no records.",
  "code": "124519405",
},
"_links": {
  "self": {
    "href": "/api/svm/svms/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/top-
metrics/files?top_metric=throughput.write"
  }
}
}
```

---

### List files with the most I/O activity

GET `/svm/svms/{svm.uuid}/top-metrics/files`

**Introduced In:** 9.11

Retrieves a list of files with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| max_records_per_svm | integer | query | False | Max records per svm. |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |
| junction-path | string | query | False | Filter by junction-path |
| svm.name | string | query | False | Filter by svm.name |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| throughput.read | integer | query | False | Filter by throughput.read |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.write | integer | query | False | Filter by throughput.write |
| path | string | query | False | Filter by path |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| excluded_volumes | array[top_metrics_svm_file_excluded_volume] | List of volumes that are not included in the SVM activity tracking REST API. |
| notice | notice | Optional field that indicates why no records are returned by the SVM activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_svm_file] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "excluded_volumes": [
    {
      "reason": {
        "code": "111411207",
        "message": "The volume is offline."
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "notice": {
    "code": "111411207",
    "message": "The volume is offline."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 5,
        "write": 4
      },
      "junction-path": "/fv",
      "path": "/vol/fv/dir_abc/dir_123/file_1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 2,
        "write": 20
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124519405 | The activity tracking report for SVM svm.name returned zero records. Check whether the activity tracking enabled volumes belonging to the SVM have read/write traffic. Refer to the REST API documentation for more information on why there might be no records. |
| 124519406 | Failed to get the activity tracking report for SVM svm.name. Reason:<Reason for="" failure="">.</Reason> |
| 124519407 | SVM wildcard queries are not supported for activity tracking reports. |
| 124519408 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination for SVM DR. |
| 124519409 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination of a MetroCluster SVM relationship and the SVM admin state is stopped. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

reason

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why the volume is not included in the SVM activity tracking REST API. |
| message | string | Details why the volume is not included in the SVM activity tracking REST API. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_svm_file_excluded_volume

List of volumes and their details as to why they are not included in the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| reason | reason | |
| volume | volume | |

notice

Optional field that indicates why no records are returned by the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

_links

| Name | Type | Description |
| --- | --- | --- |
| metadata | href | |
| self | href | |

top_metric_value_error_bounds

| Name | Type | Description |
| --- | --- | --- |
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |

| Name | Type | Description |
|---|---|---|
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

top_metrics_svm_file

Information about a file's IO activity.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| iops | iops | |
| junction-path | string | Junction path of the file. |
| path | string | Path of the file. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Retrieve users

## Retrieve users

### Overview

You can use this API to retrieve a list of users with the most I/O activity for FlexVol and FlexGroup volumes belonging to a specified SVM, within the past several seconds. To obtain this list, only the volumes which have the activity tracking feature enabled are considered.

This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` and `throughput` metrics. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API supports returning only one I/O activity type per request.

### Enabling and disabling activity tracking feature

The following APIs can be used to enable, disable, and retrieve the activity tracking state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"on"}'

– PATCH /api/storage/volumes/{uuid} -d '{"activity_tracking.state":"off"}'

– GET /api/storage/volumes/{uuid}/?fields=activity_tracking

### Excluded volumes list

Optionally, the API returns an excluded list of activity tracking-enabled volumes, which were not accounted for when obtaining the list of clients with the most I/O activity for the SVM. This excluded list contains both the volume information and the reason for exclusion.

### Approximate accounting and error bars

When too many users have recent activity, some users might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the users that would have otherwise been included might not be present in the final list returned by the API.

### Failure to return list of users with most I/O activity

The API can sometimes fail to return the list of users with the most I/O activity, due to the following reasons.

– The volumes belonging to the SVM do not have the activity tracking feature enabled.

– The volumes belonging to the SVM have not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of users with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return the usernames

The API can sometimes fail to obtain the usernames for the list of userid entries, due to internal transient errors.<br> In such cases, instead of the username, the API will return "{<user-id>}" for the user entry.</user-id>

#### Retrieve a list of the users with the most I/O activity

For a report on the users with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of the users with the greatest number of the average read operations per second. The current maximum number of users returned by the API for an I/O activity type is 25.

– GET /api/svm/svms/{svm.uuid}/top-metrics/users

#### Examples

### Retrieving a list of the users with the greatest average number of read bytes received per second:

```
# The API:
GET /api/svm/svms/{svm.uuid}/top-metrics/users

# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/users?top_metric=throughput.read"

# The Response:
{
"records": [
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
    "user_name": "user1",
    "volumes": [
      {
        "name": "fv1",
        "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
```

```
        "_links": {
          "self": {
            "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
          }
        }
      },
      {
        "name": "fv2",
        "uuid": "2ea74c3e-d5ca-11eb-8fbb-005056ac0f33",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/2ea74c3e-d5ca-11eb-8fbb-
005056ac0f33"
          }
        }
      },
      {
        "name": "fv4",
        "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
          }
        }
      }
    ]
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1002",
    "user_name": "user2",
    "volumes": [
      {
        "name": "fv2",
        "uuid": "2ea74c3e-d5ca-11eb-8fbb-005056ac0f33",
```

```
        "_links": {
          "self": {
            "href": "/api/storage/volumes/2ea74c3e-d5ca-11eb-8fbb-
005056ac0f33"
          }
        }
      },
      {
        "name": "fv3",
        "uuid": "1ca74c3e-d5ca-11eb-8fbb-005056ac0f88",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/1ca74c3e-d5ca-11eb-8fbb-
005056ac0f88"
          }
        }
      }
    ]
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1003",
    "user_name": "user3",
    "volumes": [
      {
        "name": "fv3",
        "uuid": "1ca74c3e-d5ca-11eb-8fbb-005056ac0f88",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/1ca74c3e-d5ca-11eb-8fbb-
005056ac0f88"
          }
        }
      }
    ]
  },
  {
```

```
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 235,
      "error": {
        "lower_bound": 235,
        "upper_bound": 245
      }
    },
    "user_id": "1988",
    "user_name": "user4",
    "volumes": [
      {
        "name": "fv4",
        "uuid": "5bbfc226-3fd8-42c9-a651-fa6167c2cf84",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/5bbfc226-3fd8-42c9-a651-
fa6167c2cf84"
          }
        }
      }
    ]
  },
  {
    "svm": {
      "name": "vs1"
    },
    "iops": {
      "write": 235,
      "error": {
        "lower_bound": 235,
        "upper_bound": 245
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1005",
    "user_name": "user5",
    "volumes": [
      {
        "name": "fv5",
        "uuid": "5bbfc227-3fd8-42c9-a651-fa6167c2cf85",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/5bbfc227-3fd8-42c9-a651-
fa6167c2cf85"
```

```
          }
        }
      }
    ]
  }
],
"num_records": 5,
"excluded_volumes": [
  {
    "volume": {
      "uuid": "5bbfc224-3fd8-42c9-a651-fa6167c2cf84",
      "name": "vol1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf84"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 12345
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf84"
      }
    }
  },
  {
    "volume": {
      "uuid": "5bbfc224-3fd8-42c9-a651-fa6167c2cf85",
      "name": "vol2",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf85"
        }
      }
    },
    "reason": {
      "message": "The volume is offline.",
      "code": 23456
    },
    "_links": {
```

```
      "self": {
        "href": "/api/storage/volumes/5bbfc224-3fd8-42c9-a651-
fa6167c2cf85"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/svm/svms/f5516c44-3e61-11ec-9cd5-005056acae40/top-
metrics/users?top_metric=iops.write&fields=**"
  }
}
}
```

**Example showing the behavior of the API where there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/{svm.uuid}/top-
metrics/users?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for SVM \"vs1\" returned zero
records. Check whether the activity tracking enabled volumes belonging to
the SVM have read/write traffic. Refer to the REST API documentation for
more information on why there might be no records.",
  "code": "124519405",
},
"_links": {
  "self": {
    "href": "/api/svm/svms/f5516c44-3e61-11ec-9cd5-005056acae40/top-
metrics/users?top_metric=throughput.write"
  }
}
}
```

### List users with the most I/O activity

GET `/svm/svms/{svm.uuid}/top-metrics/users`

**Introduced In:** 9.11

Retrieves a list of users with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | path | True | SVM UUID |
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| user_id | string | query | False | Filter by user_id |
| user_name | string | query | False | Filter by user_name |
| svm.name | string | query | False | Filter by svm.name |
| volumes.name | string | query | False | Filter by volumes.name<br>• Introduced in: 9.12 |
| volumes.uuid | string | query | False | Filter by volumes.uuid<br>• Introduced in: 9.12 |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| excluded_volumes | array[top_metrics_svm_user_excluded_volume] | List of volumes that are not included in the SVM activity tracking REST API. |
| notice | notice | Optional field that indicates why no records are returned by the SVM activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_svm_user] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "excluded_volumes": [
    {
      "reason": {
        "code": "111411",
        "message": "The volume is offline."
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "notice": {
    "code": "111411207",
    "message": "The volume is offline."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
```

```
        "read": 4,
        "write": 8
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 10,
        "write": 7
      },
      "user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "user_name": "James",
      "volumes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ]
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 124519405 | The activity tracking report for SVM svm.name returned zero records. Check whether the activity tracking enabled volumes belonging to the SVM have read/write traffic. Refer to the REST API documentation for more information on why there might be no records. |
| 124519406 | Failed to get the activity tracking report for SVM svm.name. Reason:<Reason for="" failure="">.</Reason> |
| 124519407 | SVM wildcard queries are not supported for activity tracking reports. |
| 124519408 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination for SVM DR. |
| 124519409 | Activity tracking is not supported on SVM svm.name, because it is configured as a destination of a MetroCluster SVM relationship and the SVM admin state is stopped. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

reason

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why the volume is not included in the SVM activity tracking REST API. |
| message | string | Details why the volume is not included in the SVM activity tracking REST API. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_svm_user_excluded_volume

List of volumes and their details as to why they are not included in the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| reason | reason | |
| volume | volume | |

notice

Optional field that indicates why no records are returned by the SVM activity tracking REST API.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
| --- | --- | --- |
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

top_metric_value_error_bounds

| Name | Type | Description |
| --- | --- | --- |
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
| --- | --- | --- |
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

volumes

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.8<br>• x-nullable: true |

top_metrics_svm_user

Aggregated information about a user's IO activity at a SVM scope.

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| user_id | string | User ID of the user. |
| user_name | string | Name of the user. |
| volumes | array[volumes] | List of volumes where the user is generating traffic. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage web services security configuration

### Manage web services security configuration

You can use this API to update and retrieve the web services security configuration for each data SVM.

#### Updating the web services security configuration

The following fields can be used to update the web services security configuration:

- certificate.uuid
- client_enabled
- ocsp_enabled
  When updating the certificate, the certificate UUID of an existing certificate known to ONTAP must be provided. The certificate must be of type "server".

A "client-ca" certificate must be installed on ONTAP to enable "client_enabled".

#### Examples

### Retrieving the web services security configuration

```
# API:
GET /api/svm/svms/{svm.uuid}/web

# The call:
curl -X GET "https://<mgmt-ip>/api/svm/svms/3c1b259d-5789-a2eb-9301-
10705682b34f/web" -H "accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "3c1b259d-5789-a2eb-9301-10705682b34f",
  "name": "svm2",
  "_links": {
    "self": {
      "href": "/api/svm/svms/3c1b259d-5789-a2eb-9301-10705682b34f"
    }
  }
},
"certificate": {
  "uuid": "a3bb219d-4382-1fe0-9c06-1070568ea23d",
  "name": "cert1",
  "_links": {
    "self": {
      "href": "/api/security/certificates/a3bb219d-4382-1fe0-9c06-
1070568ea23d"
    }
  }
},
"client_enabled": false,
"ocsp_enabled": false,
"_links": {
  "self": {
    "href": "/api/svm/svms/3c1b259d-5789-a2eb-9301-10705682b34f/web"
  }
}
}
```

**Updating the web services security configuration**

```
# The API:
PATCH /api/svm/svms/{svm.uuid}/web

# The call:
curl -X PATCH "https://<mgmt-ip>/api/svm/svms/3c1b259d-5789-a2eb-9301-
10705682b34f/web" -d '{ "certificate": {"uuid": "56da2799-13bc-2ae4-0c16-
0c71244ea2ca" }}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 28 May 2021 09:36:43 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
```

### Retrieve the web services security configuration

GET `/svm/svms/{svm.uuid}/web`

**Introduced In:** 9.10

Retrieves the web services security configuration.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |

| Name | Type | Description |
| --- | --- | --- |
| client_enabled | boolean | Indicates whether client authentication is enabled. |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Certificate name |
| uuid | string | Certificate UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the web services security configuration

PATCH `/svm/svms/{svm.uuid}/web`

**Introduced In:** 9.10

Updates the web services security configuration.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| svm.uuid | string | path | True | UUID of the SVM to which this object belongs. |

**Request Body**

| Name | Type | Description |
|---|---|---|
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |
| client_enabled | boolean | Indicates whether client authentication is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 9830442 | Client authentication cannot be enabled without a client ca certificate. |
| 9830488 | The certificate is not a "server" certificate. |
| 9830489 | The certificate does not exist for the given SVM. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

| Name | Type | Description |
|---|---|---|
| name | string | Certificate name |
| uuid | string | Certificate UUID |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

web_svm

| Name | Type | Description |
|---|---|---|
| certificate | certificate | Certificate used by cluster and node management interfaces for TLS connection requests. |
| client_enabled | boolean | Indicates whether client authentication is enabled. |
| ocsp_enabled | boolean | Indicates whether online certificate status protocol verification is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Security

## ONTAP REST API Security endpoints

### Overview

You can use ONTAP security APIs to manage security settings for the cluster and SVMs.

#### SAML

Configure the SAML 2.0 SP (Service Provider) protocol inside ONTAP. Doing so redirects the authentication task to a third-party Identity Provider (IDP) that can utilize any number of approaches for multi-factor authentication. After SAML authentication is enabled, all interactive web access (System Manager, SPI) is authenticated via SAML and a third-party IDP.

## Manage security-related operations

### Manage security-related operations

#### Overview

You can use this API for various cluster-wide security-related operations.

#### "onboard_key_manager_configurable_status" object

Use this API to retrieve details of whether or not the Onboard Key Manager can be configured on the cluster.

– GET /api/security

– GET /api/security?fields=onboard_key_manager_configurable_status

#### "software_data_encryption" object

Contains software data encryption related information.

The following APIs can be used to enable or disable and obtain default software data at rest encryption values:

– PATCH /api/security -d '{ "software_data_encryption.disabled_by_default" : true }'

– PATCH /api/security -d '{ "software_data_encryption.disabled_by_default" : false }'

– GET /api/security

– GET /api/security?fields=software_data_encryption

A PATCH request on this API using the parameter "software_data_encryption.conversion_enabled" triggers the conversion of all non-encrypted metadata volumes to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the conversion to start, the cluster must have either Onboard Key Manager or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates. For MetroCluster configurations, the PATCH request will fail if the cluster is in the switchover state.

The following API can be used to initiate software data encryption conversion.

– PATCH /api/security -d '{ "software_data_encryption.conversion_enabled" : true }'

#### "fips" object

Contains FIPS mode information.

A PATCH request on this API using the parameter "fips.enabled" switches the system from using the default cryptographic module software implementations to validated ones or vice versa, where applicable. If the value of the parameter is "true" and unapproved algorithms are configured as permitted in relevant subsystems, those algorithms will be disabled in the relevant subsystem configurations. If "false", there will be no implied change to the relevant subsystem configurations.

– GET /api/security

– GET /api/security?fields=fips

– PATCH /api/security -d '{ "fips.enabled" : true }'

– PATCH /api/security -d '{ "fips.enabled" : false }'

#### "tls" object

Contains TLS configuration information.

A PATCH request on this API using the parameter "tls.cipher_suites" and/or "tls.protocol_versions" configures

the permissible cipher suites and/or protocol versions for all TLS-enabled applications in the system. All protocol versions at or above the lowest version level specified are enabled, including those that are not explicitly specified.

– GET /api/security

– GET /api/security?fields=tls

– PATCH /api/security -d '{ "tls" : { "protocol_versions" : ["TLSv1.3", "TLSv1.2"], "cipher_suites" : ["TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384", "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256"] } }'

#### "management_protocols" object

Contains Security Protocols information.

This security protocols endpoint is used to retrieve and configure security protocols.

– GET /api/security

– GET /api/security?fields=management_protocols

– PATCH /api/security -d '{ "management_protocols" : { "rsh_enabled" : true } }'

– PATCH /api/security -d '{ "management_protocols" : { "rsh_enabled" : false } }'

– PATCH /api/security -d '{ "management_protocols" : { "telnet_enabled" : true } }'

– PATCH /api/security -d '{ "management_protocols" : { "telnet_enabled" : false } }'

– PATCH /api/security -d '{ "management_protocols" : { "rsh_enabled" : true, "telnet_enabled" : true } }'

#### GET Examples

### Retrieving information about the security configured on the cluster

The following example shows how to retrieve the configuration of the cluster.

```
# The API:
GET /api/security:

# The call:
curl -X GET 'https://<mgmt-ip>/api/security?fields=*' -H 'accept:
application/hal+json'

# The response:
{
"onboard_key_manager_configurable_status": {
  "supported": false,
  "message": "Onboard Key Manager cannot be configured on the cluster.
There are no self-encrypting disks in the cluster, and the following nodes
do not support volume granular encryption: ntap-vsim2.",
```

```
    "code": 65537300
  },
  "software_data_encryption": {
    "conversion_enabled": false,
    "disabled_by_default": false,
    "encryption_state": "unencrypted"
  },
  "fips": {
    "enabled": false
  },
  "tls": {
    "cipher_suites": [
      "TLS_RSA_WITH_AES_128_CCM",
      "TLS_RSA_WITH_AES_128_CCM_8",
      "TLS_RSA_WITH_AES_128_GCM_SHA256",
      "TLS_RSA_WITH_AES_128_CBC_SHA",
      "TLS_RSA_WITH_AES_128_CBC_SHA256",
      "TLS_RSA_WITH_AES_256_CCM",
      "TLS_RSA_WITH_AES_256_CCM_8",
      "TLS_RSA_WITH_AES_256_GCM_SHA384",
      "TLS_RSA_WITH_AES_256_CBC_SHA",
      "TLS_RSA_WITH_AES_256_CBC_SHA256",
      "TLS_RSA_WITH_ARIA_128_GCM_SHA256",
      "TLS_RSA_WITH_ARIA_256_GCM_SHA384",
      "TLS_RSA_WITH_CAMELLIA_128_CBC_SHA",
      "TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256",
      "TLS_RSA_WITH_CAMELLIA_256_CBC_SHA",
      "TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256",
      "TLS_DHE_DSS_WITH_AES_128_GCM_SHA256",
      "TLS_DHE_DSS_WITH_AES_128_CBC_SHA",
      "TLS_DHE_DSS_WITH_AES_128_CBC_SHA256",
      "TLS_DHE_DSS_WITH_AES_256_GCM_SHA384",
      "TLS_DHE_DSS_WITH_AES_256_CBC_SHA",
      "TLS_DHE_DSS_WITH_AES_256_CBC_SHA256",
      "TLS_DHE_DSS_WITH_ARIA_128_GCM_SHA256",
      "TLS_DHE_DSS_WITH_ARIA_256_GCM_SHA384",
      "TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA",
      "TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA256",
      "TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA",
      "TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA256",
      "TLS_DHE_PSK_WITH_AES_128_CBC_SHA",
      "TLS_DHE_PSK_WITH_AES_128_CBC_SHA256",
      "TLS_DHE_PSK_WITH_AES_128_CCM",
      "TLS_PSK_DHE_WITH_AES_128_CCM_8",
      "TLS_DHE_PSK_WITH_AES_128_GCM_SHA256",
      "TLS_DHE_PSK_WITH_AES_256_CBC_SHA",
```

```
    "TLS_DHE_PSK_WITH_AES_256_CBC_SHA384",
    "TLS_DHE_PSK_WITH_AES_256_CCM",
    "TLS_PSK_DHE_WITH_AES_256_CCM_8",
    "TLS_DHE_PSK_WITH_AES_256_GCM_SHA384",
    "TLS_DHE_PSK_WITH_ARIA_128_GCM_SHA256",
    "TLS_DHE_PSK_WITH_ARIA_256_GCM_SHA384",
    "TLS_DHE_PSK_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_DHE_PSK_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_DHE_PSK_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_DHE_RSA_WITH_AES_128_CCM",
    "TLS_DHE_RSA_WITH_AES_128_CCM_8",
    "TLS_DHE_RSA_WITH_AES_128_GCM_SHA256",
    "TLS_DHE_RSA_WITH_AES_128_CBC_SHA",
    "TLS_DHE_RSA_WITH_AES_128_CBC_SHA256",
    "TLS_DHE_RSA_WITH_AES_256_CCM",
    "TLS_DHE_RSA_WITH_AES_256_CCM_8",
    "TLS_DHE_RSA_WITH_AES_256_GCM_SHA384",
    "TLS_DHE_RSA_WITH_AES_256_CBC_SHA",
    "TLS_DHE_RSA_WITH_AES_256_CBC_SHA256",
    "TLS_DHE_RSA_WITH_ARIA_128_GCM_SHA256",
    "TLS_DHE_RSA_WITH_ARIA_256_GCM_SHA384",
    "TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA",
    "TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA",
    "TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256",
    "TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_ECDHE_RSA_WITH_ARIA_128_GCM_SHA256",
    "TLS_ECDHE_RSA_WITH_ARIA_256_GCM_SHA384",
    "TLS_ECDHE_ECDSA_WITH_AES_128_CCM",
    "TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8",
    "TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256",
    "TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA",
    "TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256",
    "TLS_ECDHE_ECDSA_WITH_AES_256_CCM",
    "TLS_ECDHE_ECDSA_WITH_AES_256_CCM_8",
    "TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384",
    "TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA",
    "TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384",
    "TLS_ECDHE_ECDSA_WITH_ARIA_128_GCM_SHA256",
    "TLS_ECDHE_ECDSA_WITH_ARIA_256_GCM_SHA384",
    "TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA",
    "TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA256",
    "TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA",
```

```
    "TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA384",
    "TLS_ECDHE_PSK_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_ECDHE_PSK_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_ECDHE_PSK_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256",
    "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA",
    "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256",
    "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384",
    "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA",
    "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384",
    "TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_PSK_WITH_AES_128_CBC_SHA",
    "TLS_PSK_WITH_AES_128_CBC_SHA256",
    "TLS_PSK_WITH_AES_128_CCM",
    "TLS_PSK_WITH_AES_128_CCM_8",
    "TLS_PSK_WITH_AES_128_GCM_SHA256",
    "TLS_PSK_WITH_AES_256_CBC_SHA",
    "TLS_PSK_WITH_AES_256_CBC_SHA384",
    "TLS_PSK_WITH_AES_256_CCM",
    "TLS_PSK_WITH_AES_256_CCM_8",
    "TLS_PSK_WITH_AES_256_GCM_SHA384",
    "TLS_PSK_WITH_ARIA_128_GCM_SHA256",
    "TLS_PSK_WITH_ARIA_256_GCM_SHA384",
    "TLS_PSK_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_PSK_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_PSK_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_RSA_PSK_WITH_AES_128_CBC_SHA",
    "TLS_RSA_PSK_WITH_AES_128_CBC_SHA256",
    "TLS_RSA_PSK_WITH_AES_128_GCM_SHA256",
    "TLS_RSA_PSK_WITH_AES_256_CBC_SHA",
    "TLS_RSA_PSK_WITH_AES_256_CBC_SHA384",
    "TLS_RSA_PSK_WITH_AES_256_GCM_SHA384",
    "TLS_RSA_PSK_WITH_ARIA_128_GCM_SHA256",
    "TLS_RSA_PSK_WITH_ARIA_256_GCM_SHA384",
    "TLS_RSA_PSK_WITH_CAMELLIA_128_CBC_SHA256",
    "TLS_RSA_PSK_WITH_CAMELLIA_256_CBC_SHA384",
    "TLS_RSA_PSK_WITH_CHACHA20_POLY1305_SHA256",
    "TLS_SRP_SHA_WITH_AES_128_CBC_SHA",
    "TLS_SRP_SHA_WITH_AES_256_CBC_SHA",
    "TLS_SRP_SHA_DSS_WITH_AES_128_CBC_SHA",
    "TLS_SRP_SHA_DSS_WITH_AES_256_CBC_SHA",
    "TLS_SRP_SHA_RSA_WITH_AES_128_CBC_SHA",
    "TLS_SRP_SHA_RSA_WITH_AES_256_CBC_SHA",
    "TLS_AES_128_GCM_SHA256",
```

```
    "TLS_AES_256_GCM_SHA384",
    "TLS_CHACHA20_POLY1305_SHA256"
  ],
  "protocol_versions": [
    "TLSv1.3",
    "TLSv1.2"
  ]
},
"management_protocols": {
  "rsh_enabled": false,
  "telnet_enabled": false
 }
}
```

---

#### PATCH Examples

### Enabling software encryption conversion in the cluster

The following example shows how to enable software encryption conversion in the cluster.

```
# The API:
PATCH /api/security

# The call
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{
"software_data_encryption.conversion_enabled" : true }'

# The response:
{
 "job": {
     "uuid": "ebcbd82d-1cd4-11ea-8f75-005056ac4adc",
     "_links": {
         "self": {
             "href": "/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-
005056ac4adc"
         }
     }
 }
}
```

This returns a job UUID. A subsequent GET for this job UUID returns details of the job.

```
# The API:
/api/cluster/jobs/{uuid}

# The call
curl -X GET "https://<mgmt_ip>/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-
005056ac4adc"

# The response:
{
"uuid": "ebcbd82d-1cd4-11ea-8f75-005056ac4adc",
"description": "PATCH /api/security",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2019-12-12T06:45:40-05:00",
"end_time": "2019-12-12T06:45:40-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-005056ac4adc"
  }
}
}
```

#### Enabling FIPS mode in the cluster

The following example shows how to enable FIPS mode in the cluster.

```
# The API:
PATCH /api/security

# The call
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{ "fips.enabled" : true
}'

# The response:
{
 "job": {
     "uuid": "8e7f59ee-a9c4-4faa-9513-bef689bbf2c2",
     "_links": {
         "self": {
             "href": "/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-
bef689bbf2c2"
         }
     }
 }
}
```

This returns a job UUID. A subsequent GET for this job UUID returns details of the job.

```
# The API:
/api/cluster/jobs/{uuid}

# The call
curl -X GET "https://<mgmt_ip>/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-
bef689bbf2c2"

# The response:
{
"uuid": "8e7f59ee-a9c4-4faa-9513-bef689bbf2c2",
"description": "PATCH /api/security",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-04-28T06:55:40-05:00",
"end_time": "2020-04-28T06:55:41-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-bef689bbf2c2"
  }
}
}
```

#### Configuring permissible TLS protocols and cipher suites in the cluster

The following example shows how to configure the cluster to only allow TLSv1.3 & TLSv1.2 with selected cipher suites.

```
# The API:
PATCH /api/security

# The call
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{ "tls" : {
"protocol_versions" : ["TLSv1.3", "TLSv1.2"], "cipher_suites" :
["TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384", "TLS_AES_256_GCM_SHA384"] } }'

# The response:
{
 "job": {
     "uuid": "b45b6290-f4f2-442a-aa0e-4d3ffefe5e0d",
     "_links": {
         "self": {
             "href": "/api/cluster/jobs/b45b6290-f4f2-442a-aa0e-
4d3ffefe5e0d"
         }
     }
 }
}
```

This returns a job UUID. A subsequent GET for this job UUID returns details of the job.

```
# The API:
/api/cluster/jobs/{uuid}

# The call
curl -X GET "https://<mgmt_ip>/api/cluster/jobs/b45b6290-f4f2-442a-aa0e-
4d3ffefe5e0d"

# The response:
{
"uuid": "b45b6290-f4f2-442a-aa0e-4d3ffefe5e0d",
"description": "PATCH /api/security",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-03-22T08:52:50-05:00",
"end_time": "2021-03-22T08:52:51-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/b45b6290-f4f2-442a-aa0e-4d3ffefe5e0d"
  }
}
}
```

#### Enabling security protocols in the cluster

The following example shows how to enable the security protocol rsh in the cluster.

```
# The API:
PATCH /api/security

# The call
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{
"management_protocols" : { "rsh_enabled" : true } }'

# The response
{
"job": {
"uuid": "2980ba28-adab-11eb-8fa3-005056bbfa84",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/2980ba28-adab-11eb-8fa3-005056bbfa84"
  }
 }
 }
}

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/security/?fields=management_protocols"

# The response:
{
"management_protocols": {
  "rsh_enabled": false,
  "telnet_enabled": false
},
"_links": {
  "self": {
    "href": "/api/security"
  }
}
}
```

**Retrieve information about security configured on the cluster**

GET `/security`

**Introduced In:** 9.7

Retrieves information about the security configured on the cluster.

#### Related ONTAP commands

- `security config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| fips | fips | Cluster-wide Federal Information Processing Standards (FIPS) mode information. |
| management_protocols | management_protocols | Cluster-wide security protocols related information. |
| onboard_key_manager_configurable_status | onboard_key_manager_configurable_status | Indicates whether the Onboard Key Manager can be configured in the cluster. |
| software_data_encryption | software_data_encryption | Cluster-wide software data encryption related information. |
| tls | tls | Cluster-wide Transport Layer Security (TLS) configuration information |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "onboard_key_manager_configurable_status": {
    "code": 65537300,
    "message": "No platform support for volume encryption in following
nodes - node1, node2."
  },
  "software_data_encryption": {
    "encryption_state": "string"
  },
  "tls": {
    "cipher_suites": [
      "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384"
    ],
    "protocol_versions": [
      "string"
    ]
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

fips

Cluster-wide Federal Information Processing Standards (FIPS) mode information.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the software FIPS mode is enabled on the cluster. Our FIPS compliance involves configuring the use of only approved algorithms in applicable contexts (for example TLS), as well as the use of formally validated cryptographic module software implementations, where applicable. The US government documents concerning FIPS 140-2 outline the relevant security policies in detail. |

management_protocols

Cluster-wide security protocols related information.

| Name | Type | Description |
|---|---|---|
| rsh_enabled | boolean | Indicates whether or not security protocol rsh is enabled on the cluster. |
| telnet_enabled | boolean | Indicates whether or not security protocol telnet is enabled on the cluster. |

onboard_key_manager_configurable_status

Indicates whether the Onboard Key Manager can be configured in the cluster.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Code corresponding to the status message. Returns a 0 if the Onboard Key Manager can be configured in the cluster. |
| message | string | Reason that Onboard Key Manager cannot be configured in the cluster. |
| supported | boolean | Set to true if the Onboard Key Manager can be configured in the cluster. |

software_data_encryption

Cluster-wide software data encryption related information.

| Name | Type | Description |
| --- | --- | --- |
| conversion_enabled | boolean | Indicates whether or not software encryption conversion is enabled on the cluster. A PATCH request initiates the conversion of all non-encrypted metadata volumes in the cluster to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the PATCH request to start, the cluster must have either an Onboard or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates in the cluster. For MetroCluster configurations, a PATCH request enables conversion on all the aggregates and metadata volumes of both local and remote clusters and is not allowed when the MetroCluster is in switchover state. |
| disabled_by_default | boolean | Indicates whether or not default software data at rest encryption is disabled on the cluster. |

| Name | Type | Description |
| --- | --- | --- |
| encryption_state | string | Software data encryption state.<br> encrypted ‐ All the volumes are encrypted.<br> encrypting ‐ Encryption conversion operation is in progress.<br> partial ‐ Some volumes are encrypted, and others remains in plain text.<br> rekeying ‐ All volumes are currently being encrypted with a new key.<br> unencrypted ‐ None of the volumes are encrypted.<br> conversion_paused ‐ Encryption conversion operation is paused on one or more volumes.<br> rekey_paused ‐ Encryption rekey operation is paused on one or more volumes. |
| rekey | boolean | |

tls

Cluster-wide Transport Layer Security (TLS) configuration information

| Name | Type | Description |
| --- | --- | --- |
| cipher_suites | array[string] | Names a cipher suite that the system can select during TLS handshakes. A list of available options can be found on the Internet Assigned Number Authority (IANA) website. |
| protocol_versions | array[string] | Names a TLS protocol version that the system can select during TLS handshakes. The use of SSLv3 or TLSv1 is discouraged. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the software FIPS mode or enable conversion of non-encrypted metadata volumes non-NAE aggregates

PATCH `/security`

**Introduced In:** 9.8

Updates the software FIPS mode or modifies software data encryption.
The PATCH request can be used to enable conversion of non-encrypted metadata volumes to encrypted metadata volumes and non-NAE aggregates to NAE aggregates.

#### Related ONTAP commands

- `security config modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| fips | fips | Cluster-wide Federal Information Processing Standards (FIPS) mode information. |
| management_protocols | management_protocols | Cluster-wide security protocols related information. |
| software_data_encryption | software_data_encryption | Cluster-wide software data encryption related information. |

| Name | Type | Description |
| --- | --- | --- |
| tls | tls | Cluster-wide Transport Layer Security (TLS) configuration information |

**Example request**

```
{
  "software_data_encryption": {
    "encryption_state": "string"
  },
  "tls": {
    "cipher_suites": [
      "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384"
    ],
    "protocol_versions": [
      "string"
    ]
  }
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| job | job_link | |

**Example response**

```
{
  "job": {
    "uuid": "string"
  }
}
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5636142 | This operation is not supported in a mixed-release cluster. |
| 5636145 | This operation is not supported when cluster security is configured with FIPS mode. |
| 52428817 | SSLv3 is not supported when FIPS is enabled. |
| 52428824 | TLSv1 is not supported when FIPS is enabled. |
| 52428830 | Cannot enable FIPS-compliant mode because the configured minimum security strength for certificates is not compatible. |
| 52428832 | TLSv1.1 is not supported when FIPS is enabled. |
| 52559974 | Cannot enable FIPS-compliant mode because a certificate that is not FIPS-compliant is in use. |
| 65536987 | One or more key servers are unavailable. |
| 196608047 | Operation is not allowed when volume move is in progress. |
| 196608070 | Key manager is not configured on the cluster. Configure either an external Key Management Server or an onboard key manager. |
| 196608081 | Cannot start software encryption conversion while there are data volumes in the cluster. |
| 196608082 | The operation is not valid when the MetroCluster is in switchover mode. |
| 196608368 | Failed to perform the requested operation. One or more data volume in offline state. |
| 196608369 | Conversion cannot be enabled because the cluster contains read-only or primordial logical data-protection volumes. Retry the patch operation after deleting those volumes. |
| 196608370 | The conversion of non-encrypted volumes to NVE (NetApp Volume Encryption) volumes is already running. Monitor the NVE conversion status by querying the "software_data_encryption.encryption_state" field on the "/api/security" endpoint. |

| Error Code | Description |
| --- | --- |
| 196608371 | NVE (NetApp Volume Encryption) volumes are already being rekeyed. Monitor the NVE rekey status by querying the "software_data_encryption.encryption_state" field on the "/api/security" endpoint. |
| 196608372 | An automated ONTAP update is in progress, retry the PATCH request after it is completed. |
| 196608373 | Unable to perform the encryption operation because of a mixed-release cluster. Complete the upgrade or revert operation, then try the PATCH request again. |
| 196608374 | Failed to perform the requested operation. One or more SVMs not in admin running state. |
| 196608375 | Failed to perform the requested operation. One or more volume is of temporary type. |
| 196608376 | Internal error. Could not get volume encryption information. |
| 196608377 | Internal error. The Volume Location Database (VLDB) is inconsistent. Contact support personnel to resolve this issue. |
| 196608378 | Failed to perform the requested operation. Data SVM Key manager configuration is in mixed state. |
| 196608379 | Internal error. The encryption metadata for the volume is inconsistent. Contact technical support for assistance. |
| 196608380 | Failed to perform the requested operation. Wafliron is currently active. |
| 196608381 | Failed to perform the requested operation. A clone split operation is in progress. |
| 196608382 | Failed to perform the requested operation. A volume rehost operation is in progress. |
| 196608383 | Failed to perform the requested operation. The cluster contains one or more SnapLock volume. |
| 196608384 | The PATCH request to start rekey failed because the cluster contains one or more plain text volumes. Retry the PATCH request after converting the existing plain text volumes to encrypted volumes. |
| 196608385 | Failed to perform the requested operation. Keystore configuration is being switched. Wait until the keystore is in the active state and then try the PATCH request again. |

| Error Code | Description |
| --- | --- |
| 196608386 | Failed to perform the requested operation. Rekey operation for one or more SVMs is in progress. Wait until the keystore is in the active state and then try the PATCH request again. |
| 196608387 | "software_data_encryption.conversion_enabled" cannot be set to "false" in a PATCH request. |
| 196608388 | Both "software_data_encryption.conversion_enabled" and "software_data_encryption.disabled_by_default" cannot be set to "true" in a single PATCH request. |
| 196608389 | Both "software_data_encryption.conversion_enabled" and "software_data_encryption.rekey" cannot be set to "true" in a single PATCH request. |
| 196608390 | "software_data_encryption.rekey" cannot be set to "false" in a PATCH request. |
| 196608391 | Both "software_data_encryption.rekey" and "software_data_encryption.disabled_by_default" cannot be set to "true" in a single PATCH request. |
| 196608392 | The PATCH request for cluster level rekey requires an effective cluster version of 9.16.1 or later. |
| 196608393 | The PATCH request for cluster level rekey is not supported on this platform. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

fips

Cluster-wide Federal Information Processing Standards (FIPS) mode information.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether or not the software FIPS mode is enabled on the cluster. Our FIPS compliance involves configuring the use of only approved algorithms in applicable contexts (for example TLS), as well as the use of formally validated cryptographic module software implementations, where applicable. The US government documents concerning FIPS 140-2 outline the relevant security policies in detail. |

management_protocols

Cluster-wide security protocols related information.

| Name | Type | Description |
|---|---|---|
| rsh_enabled | boolean | Indicates whether or not security protocol rsh is enabled on the cluster. |
| telnet_enabled | boolean | Indicates whether or not security protocol telnet is enabled on the cluster. |

onboard_key_manager_configurable_status

Indicates whether the Onboard Key Manager can be configured in the cluster.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Code corresponding to the status message. Returns a 0 if the Onboard Key Manager can be configured in the cluster. |
| message | string | Reason that Onboard Key Manager cannot be configured in the cluster. |
| supported | boolean | Set to true if the Onboard Key Manager can be configured in the cluster. |

software_data_encryption

Cluster-wide software data encryption related information.

| Name | Type | Description |
| --- | --- | --- |
| conversion_enabled | boolean | Indicates whether or not software encryption conversion is enabled on the cluster. A PATCH request initiates the conversion of all non-encrypted metadata volumes in the cluster to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the PATCH request to start, the cluster must have either an Onboard or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates in the cluster. For MetroCluster configurations, a PATCH request enables conversion on all the aggregates and metadata volumes of both local and remote clusters and is not allowed when the MetroCluster is in switchover state. |
| disabled_by_default | boolean | Indicates whether or not default software data at rest encryption is disabled on the cluster. |

| Name | Type | Description |
| --- | --- | --- |
| encryption_state | string | Software data encryption state.<br> encrypted ‐ All the volumes are encrypted.<br> encrypting ‐ Encryption conversion operation is in progress.<br> partial ‐ Some volumes are encrypted, and others remains in plain text.<br> rekeying ‐ All volumes are currently being encrypted with a new key.<br> unencrypted ‐ None of the volumes are encrypted.<br> conversion_paused ‐ Encryption conversion operation is paused on one or more volumes.<br> rekey_paused ‐ Encryption rekey operation is paused on one or more volumes. |
| rekey | boolean | |

tls

Cluster-wide Transport Layer Security (TLS) configuration information

| Name | Type | Description |
| --- | --- | --- |
| cipher_suites | array[string] | Names a cipher suite that the system can select during TLS handshakes. A list of available options can be found on the Internet Assigned Number Authority (IANA) website. |
| protocol_versions | array[string] | Names a TLS protocol version that the system can select during TLS handshakes. The use of SSLv3 or TLSv1 is discouraged. |

security_config

| Name | Type | Description |
| --- | --- | --- |
| fips | fips | Cluster-wide Federal Information Processing Standards (FIPS) mode information. |

| Name | Type | Description |
| --- | --- | --- |
| management_protocols | management_protocols | Cluster-wide security protocols related information. |
| software_data_encryption | software_data_encryption | Cluster-wide software data encryption related information. |
| tls | tls | Cluster-wide Transport Layer Security (TLS) configuration information |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage security-related accounts

### Manage security-related accounts

#### Overview

A valid user account is required to login to and provision, monitor, and manage the cluster. The scope of the management operation can be at the cluster level or at an individual SVM level. There is a need to create user accounts with specific privileges apart from the default user accounts, "admin", for cluster and "vsadmin" for SVM. Custom user accounts can be configured to perform specific (scoped) operations. User accounts can either be created locally (on the Netapp system) or referenced from an external directory server (NIS, LDAP, or Active Directory).
Apart from creation, modification, and deletion of a user account, locking and unlocking of a user account or resetting the password (for local accounts only) is possible.

A user account must be associated with the following before it can become operational:

1. A management application (SSH, HTTP, console, service_processor, and such like) for user login. HTTP enables REST API access.
2. Scope - either cluster or SVM.
3. Authentication source - password (local, NIS/LDAP, Active Directory), public/private key pair-based, certificate based.
4. RBAC role - determines what operations are permitted for the user account.

### Restrictions

A number of internal/restricted account names, such as admin, diag, autosupport, and root cannot be used.

There must be at least one console cluster administrator account. Any attempt to delete the last remaining administrator account fails.

Multi-factor authentication is only possible for SSH and Service Processor (SP) applications. For the SSH application, the only possible combinations are password (local or NIS/LDAP/Active Directory) and public key and password or public key (local) and TOTP. For the SP application, the only valid combination is password and public key.

If an account is created using the HTTP application type, the password authentication method, and its username contains one of the following special characters: @ or \|, then you cannot use this account for basic authorization and its only purpose is for OAuth 2.0 validation.

All authentication sources are not supported by all applications. You must select a compatible authentication method based on the application. The following types of authentications methods are supported:

| Application | Supported Authentication Methods |
|---|---|
| amqp | password |
| console | password |
| service_processor | password, publickey |
| HTTP | password, domain, nsswitch, certificate |
| ONTAPI | password, domain, nsswitch, certificate |
| SSH | password, publickey (key pair), domain, nsswitch, totp |

> In this table, "totp" means time-based one-time password and is only allowed to be configured as second authentication, "certificate" means security certificate, "domain" means that the user directory server is an external Active Directory, "nsswitch" means the directory server is an external NIS or LDAP server. At login time, the user is authenticated with these external directory servers which must be provisioned separately. For the SP (*service_processor* in above table) application, *publickey* can only be configured as the second factor of authentication.

Support for publickey authentication and MFA for Domain users has been added.

Support for TOTP as a secondary authentication method with password or public key as the primary authentication method has been added.

For SP (*service_processor* in above table) application, support for *publickey* as a secondary authentication method with password as the primary authentication method has been added.

#### Examples

### Creating a cluster-scoped user account

Specify the user account name, role name, and the tuples (of application and authentication methods) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.

> Each entry in the applications array must be for a different application.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with applications ssh, http
and password authentication scheme:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"cluster_user1","applications":[{"application":"ssh","authenticat
ion_methods":["password"],"second_authentication_method":"none"},{"applica
tion":"http","authentication_methods":["password"]}],"role":"admin","passw
ord":"p@ssw@rd123"}'
Note: The password is an optional parameter for creation and can be set
later using a PATCH request. See the examples for modification of user
account or password.
```

### Creating an SVM-scoped user account

For an SVM-scoped account, specify either the SVM name as the owner.name or SVM uuid as the owner.uuid along with other parameters for the user account. These indicate the SVM for which the user account is being created and can be obtained from the response body of GET performed on the */api/svm/svms* API.

```
# The API:
POST "/api/security/accounts"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"owner":{"uuid":"aaef7c38-4bd3-11e9-b238-
0050568e2e25"},"name":"svm_user1","applications":[{"application":"ssh","au
thentication_methods":["password"],"second_authentication_method":"none"}]
,"role":"vsadmin","password":"p@ssw@rd123"}'
```

#### Retrieving the configured user accounts

Use the following API to retrieve all of the user accounts or a filtered list of user accounts (by name, for a specific SVM, and so on).

```
# The API:
GET "/api/security/accounts"

# The call to retrieve all the user accounts configured in the cluster:
curl -X GET "https://<mgmt-ip>/api/security/accounts"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "admin",
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-
0050568e2e25/admin"
      }
    }
  },
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
```

```
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "autosupport",
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-
0050568e2e25/autosupport"
            }
    }
  },
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "cluster_user1",
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-
0050568e2e25/cluster_user1"
      }
    }
  },
  {
    "owner": {
      "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "svm_user1",
    "_links": {
      "self": {
```

```
        "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_user1"
      }
    }
  },
  {
    "owner": {
      "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "vsadmin",
    "_links": {
      "self": {
        "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/security/accounts"
  }
}
}

# The scoped call to retrieve the configured cluster-scoped user accounts:
curl -X GET "https://<mgmt-ip>/api/security/accounts/?scope=cluster"

# The scoped call to retrieve the configured SVM-scoped user accounts:
curl -X GET "https://<mgmt-ip>/api/security/accounts/?scope=svm"

# The scoped call to retrieve the user accounts configured for the SVM
"svm1":
curl -X GET "https://<mgmt-ip>/api/security/accounts/?owner.name=svm1"

# The scoped call to retrieve the user accounts configured with the
"admin" role:
curl -X GET "https://<mgmt-ip>/api/security/accounts/?role=admin"
```

#### Creating an Active Directory users with publickey authentication

Specify the Active Directory user account name, role name, and the tuples (application and authentication methods) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with application ssh and
publickey authentication scheme for domain users:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"domain_name\\cluster_user__u1","applications":[{"application":"s
sh","authentication_methods":["publickey"]}]}'
```

#### Creating an Active Directory user with MFA(domain+publickey)

Specify the Active Directory user account name, role name, and the tuples (application and authentication methods) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with application ssh and and
MFA for domain users:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"domain_name\\cluster_user_u1","applications":[{"application":"ss
h","authentication_methods":["domain"],"second_authentication_method":"pub
lickey"}]}'
```

#### Retrieving the configured Active directory user accounts

Use the following API to retrieve all of the Active directory user accounts.

```
# The API:
curl -X GET "https://<mgmt-ip>/api/security/accounts/?name=*\*"

# The response:
{
"records": [
{
  "owner": {
    "uuid": "d6a740a0-4086-11ed-9f68-0050568edfd7",
    "name": "cluster-1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/d6a740a0-4086-11ed-9f68-0050568edfd7"
      }
    }
  },
  "name": "domain\\ad_user_u1",
  "_links": {
    "self": {
      "href": "/api/security/accounts/d6a740a0-4086-11ed-9f68-
0050568edfd7/domain%5Cad_user_u1"
    }
  }
}
 ],
 "num_records": 1,
 "_links": {
"self": {
  "href": "/api/security/accounts/?name=*\\*"
}
 }
}
```

**Creating a user with MFA (password+TOTP)**

Cluster-scoped user account:
Follow the cluster-scoped user creation example and additionally specify the 'totp' as the second_authentication_method.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with application ssh,
authentication password and totp:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"cluster_user_1","applications":[{"application":"ssh","authentica
tion_methods":["password"],"second_authentication_method":"totp"}]}'
```

SVM-scoped user account:
Follow the SVM-scoped user creation example and additionally specify the 'totp' as the second_authentication_method.

```
#The API

# The call to create a SVM-scoped user account with application ssh,
authentication password and totp:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"owner":{"uuid":"aaef7c38-4bd3-11e9-b238-
0050568e2e25"},"name":"svm_user1","applications":[{"application":"ssh","au
thentication_methods":["password"],"second_authentication_method":"totp"}]
,"role":"vsadmin","password":"<SVM-USER-PASSWORD>"}'
```

#### Creating an SP user with MFA (password + publickey)

Specify the SP user account name, role name, and the tuples (application and authentication methods) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with application
service_processor, authentication method password and second
authentication publickey
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"cluster_sp_user1","applications":[{"application":"service_proces
sor","authentication_methods":["password"],"second_authentication_method":
"publickey"}]}'
```

#### Retrieving the configured SP user accounts

Use the following API to retrieve all of the configured SP user accounts.

```
# The API:
curl -X GET "https://<mgmt-
ip>/api/security/accounts/?applications.application=service_processor"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "cluster_sp_user1",
    "applications": [
      {
        "application": "service_processor"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-
0050568e2e25/"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/security/accounts/?applications.application=service_processor"
  }
}
}
```

**Retrieve user accounts in the cluster**

GET `/security/accounts`

**Introduced In:** 9.6

Retrieves a list of user accounts in the cluster.

#### Related ONTAP commands

- `security login show`

#### Learn more

- DOC /security/accounts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| password_hash_algorithm | string | query | False | Filter by password_hash_algorithm<br>• Introduced in: 9.11 |
| owner.name | string | query | False | Filter by owner.name<br>• Introduced in: 9.7 |
| owner.uuid | string | query | False | Filter by owner.uuid<br>• Introduced in: 9.7 |
| locked | boolean | query | False | Filter by locked<br>• Introduced in: 9.7 |
| comment | string | query | False | Filter by comment<br>• Introduced in: 9.7 |
| scope | string | query | False | Filter by scope<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| applications.is_ns_switch_group | boolean | query | False | Filter by applications.is_ns_switch_group<br>• Introduced in: 9.15 |
| applications.second_authentication_method | string | query | False | Filter by applications.second_authentication_method<br>• Introduced in: 9.7 |
| applications.application | string | query | False | Filter by applications.application<br>• Introduced in: 9.7 |
| applications.authentication_methods | string | query | False | Filter by applications.authentication_methods<br>• Introduced in: 9.7 |
| applications.is_ldap_fastbind | boolean | query | False | Filter by applications.is_ldap_fastbind<br>• Introduced in: 9.14 |
| role.name | string | query | False | Filter by role.name<br>• Introduced in: 9.7 |
| name | string | query | False | Filter by name<br>• Introduced in: 9.7<br>• maxLength: 64<br>• minLength: 3 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[account] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "applications": [
        {
          "application": "string",
          "authentication_methods": [
            "string"
          ],
          "second_authentication_method": "string"
        }
      ],
      "comment": "string",
      "name": "joe.smith",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "password_hash_algorithm": "sha512",
      "role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
          "name": "admin"
        },
        "scope": "string"
      }
    ]
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

account_application

| Name | Type | Description |
|---|---|---|
| application | string | Applications |
| authentication_methods | array[string] | |
| is_ldap_fastbind | boolean | Optional property that specifies the mode of authentication as LDAP Fastbind. |
| is_ns_switch_group | boolean | Optional property that specifies whether the user is an LDAP or NIS group. |

| Name | Type | Description |
| --- | --- | --- |
| second_authentication_method | string | An optional additional authentication method for multi-factor authentication (MFA). This property is only supported for SSH (*ssh*) and Service Processor (*service_processor*) applications. It is ignored for all other applications. Time-based One-Time Passwords (TOTPs) are only supported with the authentication method password or public key. For the Service Processor (*service_processor*) application, *none* and *publickey* are the only supported enum values. |

owner

Owner name and UUID that uniquely identifies the user account.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

role

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | Role name |

account

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| applications | array[account_application] | |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| name | string | User or group account name |
| owner | owner | Owner name and UUID that uniquely identifies the user account. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a new user account

POST `/security/accounts`

**Introduced In:** 9.6

Creates a new user account.

#### Required parameters

- `name` - Account name to be created.
- `applications` - Array of one or more application tuples (of application and authentication methods).

#### Optional parameters

- `owner.name` or `owner.uuid` - Name or UUID of the SVM for an SVM-scoped user account. If not supplied, a cluster-scoped user account is created.
- `role` - RBAC role for the user account. Defaulted to `admin` for cluster user account and to `vsadmin` for SVM-scoped account.
- `password` - Password for the user account (if the authentication method is opted as password for one or more of applications).
- `second_authentication_method` - Needed for MFA and only supported for *ssh* and *service_processor* applications. Defaults to *none* if not supplied.
- `comment` - Comment for the user account (e.g purpose of this account).
- `locked` - Locks the account after creation. Defaults to `false` if not supplied.
- `is_ldap_fastbind` - Needed for LDAP Fastbind Authentication and only supported for applications SSH, ONTAPI, and HTTP with authentication method "nsswitch" only. Defaults to false if not supplied.
- `is_ns_switch_group` - Specifies whether the user is an LDAP or NIS group and is only supported for SSH, ONTAPI, and HTTP applications with the authentication method "nsswitch". Defaults to false if not supplied.

#### Related ONTAP commands

- `security login create`

#### Learn more

- DOC /security/accounts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| applications | array[account_application] | |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| name | string | User or group account name |
| owner | owner | Owner name and UUID that uniquely identifies the user account. |
| password | string | Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

#### Example request

```
{
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "name": "joe.smith",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "password": "string",
  "password_hash_algorithm": "sha512",
  "role": {
    "name": "admin"
  },
  "scope": "string"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1261215 | The role was not found. |
| 1261225 | Invalid command directory name. |
| 1263343 | Cannot lock user with password not set or non-password authentication method. |
| 2621475 | This operation is not supported on a node SVM. |
| 2621601 | This operation is not supported on a system SVM. |
| 2621706 | The specified owner.uuid and owner.name refer to different SVMs. |
| 5636099 | User creation with a non-admin role is not supported for service-processor application. |
| 5636121 | The user account name is reserved for use by the system. |
| 5636126 | Cannot create a user with the username or role as autosupport because it is reserved by the system. |
| 5636136 | Specifying "is_ns_switch_group" as "true" is supported only for authentication method "nsswitch". |
| 5636140 | Creating a login with application console for a data SVM is not supported. |
| 5636141 | Creating a login with application service-processor for a data SVM is not supported. |
| 5636164 | The authentication method and the second authentication method cannot be any combination of 'nsswitch' and 'password'. |
| 5636165 | Second authentication method is not supported for NIS or LDAP group based accounts. |
| 5636176 | The application and authentication-method combination is invalid. |
| 5636178 | An invalid value is specified for field "application". |
| 5636179 | Creating an AMQP application login for a data SVM is not supported. |
| 5636197 | LDAP fastbind combination for application and authentication method is not supported. |
| 5636198 | LDAP fastbind authentication is supported only for nsswitch. |
| 5636206 | Non-domain user cannot have a backslash in the username. |
| 5636207 | If the value for either the authentication-method or second-authentication-method parameters is domain, the other parameter must be publickey or none. |

| Error Code | Description |
| --- | --- |
| 5636212 | TOTP is supported only when the primary authentication method is password or public key. |
| 5636214 | Configuring the user with TOTP as a secondary authentication method requires an effective cluster version of 9.13.1 or later. |
| 5636223 | Specifying "is_ns_switch_group" as "true" is supported only for SSH, ONTAPI, and HTTP applications. |
| 5636224 | Configuring a Service Processor (SP) user with two-factor authentication requires an effective cluster version of 9.15.1 or later. |
| 5636225 | For a Service Processor (SP) user, the second factor of authentication must be one of publickey or none. |
| 5636226 | Internal error. Failed to check for ONTAP capability. |
| 5636250 | The "second_authentication_method" parameter is supported for HTTP, SSH, and Service Processor (SP) applications only. |
| 5636251 | For HTTP application, "publickey" is the only supported value for the "second_authentication_method" parameter; and can be specified only if the "authentication_methods" parameter has value "password" or "domain" or "nsswitch". |
| 7077897 | Invalid character in username. |
| 7077898 | The username must contain both letters and numbers. |
| 7077899 | The username does not meet length requirements. |
| 7077906 | A role with that name has not been defined for the Vserver. |
| 7077918 | The password cannot contain the username. |
| 7077919 | The minimum length for new password does not meet the policy. |
| 7077920 | A new password must have both letters and numbers. |
| 7077921 | The minimum number of special characters required do not meet the policy. |
| 7077929 | Cannot lock user with password not set or non-password authentication method. |
| 7077940 | The password exceeds the maximum supported length. |
| 7077941 | The defined password composition exceeds the maximum password length of 128 characters. |

| Error Code | Description |
| --- | --- |
| 7078900 | An admin password is not set. Set the password by including it in the request. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

account_application

| Name | Type | Description |
|---|---|---|
| application | string | Applications |
| authentication_methods | array[string] | |
| is_ldap_fastbind | boolean | Optional property that specifies the mode of authentication as LDAP Fastbind. |
| is_ns_switch_group | boolean | Optional property that specifies whether the user is an LDAP or NIS group. |
| second_authentication_method | string | An optional additional authentication method for multi-factor authentication (MFA). This property is only supported for SSH (*ssh*) and Service Processor (*service_processor*) applications. It is ignored for all other applications. Time-based One-Time Passwords (TOTPs) are only supported with the authentication method password or public key. For the Service Processor (*service_processor*) application, *none* and *publickey* are the only supported enum values. |

owner

Owner name and UUID that uniquely identifies the user account.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

role

| Name | Type | Description |
| --- | --- | --- |
| name | string | Role name |

account

| Name | Type | Description |
| --- | --- | --- |
| applications | array[account_application] | |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| name | string | User or group account name |
| owner | owner | Owner name and UUID that uniquely identifies the user account. |
| password | string | Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage scoped user accounts

### Manage scoped user accounts

#### Overview

This API displays and manages the configuration of scoped user accounts.

Newly created user accounts might need to be updated for many reasons. For example, a user account might need to use a different application or its role might need to be modified. According to a policy, the password or authentication source of a user account might need to be changed, or a user account might need to be locked or deleted from the system. This API allows you to make these changes to user accounts.

Specify the owner UUID and the user account name in the URI path. The owner UUID corresponds to the UUID of the SVM for which the user account has been created and can be obtained from the response body of the GET request performed on one of the following APIs:<br> /api/security/accounts for all user accounts<br> /api/security/accounts/?scope=cluster for cluster-scoped user accounts<br> /api/security/accounts/?scope=svm for SVM-scoped accounts<br> /api/security/accounts/?owner.name=<svm-name></i> for a specific SVM This API response contains the complete URI for each user account that can be used.</svm-name>

#### Examples

### Retrieving the user account details

```
# The API:
GET "/api/security/accounts/{owner.uuid}/{name}"
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/security/accounts/aef7c38-4bd3-11e9-
b238-0050568e2e25/svm_user1"

# The response:
{
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
    }
        }
},
"name": "svm_user1",
"applications": [
  {
    "application": "ssh",
    "authentication_methods": [
      "password"
    ],
    "second_authentication_method": "none"
  }
],
"role": {
  "name": "vsadmin",
  "_links": {
    "self": {
      "href": "/api/svms/aaef7c38-4bd3-11e9-b238-
0050568e2e25/admin/roles/vsadmin"
    }
  }
},
"locked": false,
"password_hash_algorithm": "sha512",
"scope": "svm",
"_links": {
  "self": {
    "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_user1"
  }
}
}
```

#### Updating the applications and role in a user account

Specify the desired configuration in the form of tuples (of applications and authentication methods) and the role. All other previously configured applications that are not specified in the "applications" parameter of the PATCH request will be de-provisioned for the user account.

```
# The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"

# The call to update the applications and role:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_user1" -d
'{"applications":[{"application":"http","authentication_methods":["domain"
]},{"application":"ontapi","authentication_methods":["password"]}],"role":
{"name": "vsadmin-backup"}}'

# The call to update only the role:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_user1" -d '{"role":"vsadmin-protocol"}'
```

#### Updating the password for a user account

```
# The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_user1" -d '{"password":"newp@ssw@rd2"}'
```

#### Locking a user account

```
The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"
The call:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_user1" -d '{"locked":"true"}'
```

#### Deleting a user account

```
# The API:
DELETE "/api/security/accounts/{owner.uuid}/{name}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-
11e9-b238-0050568e2e25/svm_user1"
```

### Delete a user account

DELETE `/security/accounts/{owner.uuid}/{name}`

**Introduced In:** 9.6

Deletes a user account.

#### Required parameters

- `name` - Account name to be deleted.
- `owner.uuid` - UUID of the SVM housing the user account to be deleted.

#### Related ONTAP commands

- `security login delete`

#### Learn more

- DOC /security/accounts/{owner.uuid}/{name}
- DOC /security/accounts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | Account owner UUID |
| name | string | path | True | User account name |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 5636098 | Last unlocked account that has an admin role cannot be deleted. |
| 5636125 | The operation is not supported on system accounts. |
| 5636146 | Cannot delete the last console account with admin role. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a specific user account

GET `/security/accounts/{owner.uuid}/{name}`

**Introduced In:** 9.6

Retrieves a specific user account.

#### Related ONTAP commands

- `security login show`

#### Learn more

- DOC /security/accounts/{owner.uuid}/{name}
- DOC /security/accounts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | Account owner UUID |
| name | string | path | True | User account name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| applications | array[account_application] | |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| name | string | User or group account name |
| owner | owner | Owner name and UUID that uniquely identifies the user account. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "name": "joe.smith",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "password_hash_algorithm": "sha512",
  "role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "scope": "string"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

account_application

| Name | Type | Description |
|---|---|---|
| application | string | Applications |
| authentication_methods | array[string] | |
| is_ldap_fastbind | boolean | Optional property that specifies the mode of authentication as LDAP Fastbind. |
| is_ns_switch_group | boolean | Optional property that specifies whether the user is an LDAP or NIS group. |
| second_authentication_method | string | An optional additional authentication method for multi-factor authentication (MFA). This property is only supported for SSH (*ssh*) and Service Processor (*service_processor*) applications. It is ignored for all other applications. Time-based One-Time Passwords (TOTPs) are only supported with the authentication method password or public key. For the Service Processor (*service_processor*) application, *none* and *publickey* are the only supported enum values. |

owner

Owner name and UUID that uniquely identifies the user account.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

role

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update a user account**

PATCH `/security/accounts/{owner.uuid}/{name}`

**Introduced In:** 9.6

Updates a user account. Locks or unlocks a user account and/or updates the role, applications, and/or password for the user account.

#### Required parameters

- `name` - Account name to be updated.
- `owner.uuid` - UUID of the SVM housing the user account to be updated.

#### Optional parameters

- `applications` - Array of one or more tuples (of application and authentication methods).
- `role` - RBAC role for the user account.
- `password` - Password for the user account (if the authentication method is opted as password for one or more of applications).
- `second_authentication_method` - Needed for MFA and only supported for *ssh* and *service_processor* applications. Defaults to *none* if not supplied.
- `comment` - Comment for the user account (e.g purpose of this account).
- `locked` - Set to true/false to lock/unlock the account.
- `is_ldap_fastbind` - Set to true/false to enable LDAP Fastbind Authentication.
- `is_ns_switch_group` - Set to true/false to specify whether the user is an LDAP or NIS group.

#### Related ONTAP commands

- `security login create`
- `security login modify`
- `security login password`
- `security login lock`
- `security login unlock`

#### Learn more

- DOC /security/accounts/{owner.uuid}/{name}
- DOC /security/accounts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.uuid | string | path | True | Account owner UUID |
| name | string | path | True | User account name |

#### Request Body

| Name | Type | Description |
|---|---|---|
| applications | array[account_application] | |

| Name | Type | Description |
| --- | --- | --- |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| password | string | Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

#### Example request

```
{
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "password": "string",
  "password_hash_algorithm": "sha512",
  "role": {
    "name": "admin"
  },
  "scope": "string"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1261215 | The role was not found. |
| 1261218 | The user was not found. |
| 1263343 | Cannot lock user with password not set or non-password authentication method. |
| 5636096 | Cannot perform the operation for this user account since the password is not set. |
| 5636097 | The operation for user account failed since user password is not set. |

| Error Code | Description |
| --- | --- |
| 5636100 | Modification of a service-processor user’s role to a non-admin role is not supported. |
| 5636125 | The operation not supported on autosupport user account which is reserved. |
| 5636129 | The role does not exist. |
| 5636136 | Specifying "is_ns_switch_group" as "true" is supported only for authentication method "nsswitch". |
| 5636154 | For mixed-version cluster the second authentication method parameter is supported for SSH and Service Processor (SP) applications only. |
| 5636159 | For a given user and application, if the second-authentication-method is specified, only one such login entry is supported. |
| 5636164 | The authentication method and the second authentication method cannot be any combination of 'nsswitch' and 'password'. |
| 5636165 | Second authentication method is not supported for NIS or LDAP group based accounts. |
| 5636197 | LDAP fastbind combination for application and authentication method is not supported. |
| 5636198 | LDAP fastbind authentication is supported only for nsswitch. |
| 5636212 | TOTP is supported only when the primary authentication method is password or public key. |
| 5636214 | Configuring the user with TOTP as a secondary authentication method requires an effective cluster version of 9.13.1 or later. |
| 5636223 | Specifying "is_ns_switch_group" as "true" is supported only for SSH, ONTAPI and HTTP applications. |
| 5636224 | Configuring a Service Processor (SP) user with two-factor authentication requires an effective cluster version of 9.15.1 or later. |
| 5636225 | For a Service Processor (SP) user, the second factor of authentication must be one of publickey or none. |
| 5636226 | Internal error. Failed to check for ONTAP capability. |
| 5636233 | User modification failed because LDAP is not configured for the SVM or the LDAP connection is not secure. |
| 7077896 | Cannot lock the account of the last console admin user. |

| Error Code | Description |
| --- | --- |
| 7077906 | A role with that name has not been defined for the Vserver. |
| 7077911 | The user is not configured to use the password authentication method. |
| 7077918 | The password cannot contain the username. |
| 7077919 | The minimum length for new password does not meet the policy. |
| 7077920 | The new password must have both letters and numbers. |
| 7077921 | The minimum number of special characters required do not meet the policy. |
| 7077924 | The new password must be different than last N passwords. |
| 7077925 | The new password must be different to the old password. |
| 7077929 | Cannot lock user with password not set or non-password authentication method. |
| 7077940 | The password exceeds maximum supported length. |
| 7077941 | Defined password composition exceeds the maximum password length of 128 characters. |
| 7078900 | An aAdmin password is not set. Set the password by including it in the request. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

account_application

| Name | Type | Description |
|---|---|---|
| application | string | Applications |
| authentication_methods | array[string] | |
| is_ldap_fastbind | boolean | Optional property that specifies the mode of authentication as LDAP Fastbind. |
| is_ns_switch_group | boolean | Optional property that specifies whether the user is an LDAP or NIS group. |
| second_authentication_method | string | An optional additional authentication method for multi-factor authentication (MFA). This property is only supported for SSH (*ssh*) and Service Processor (*service_processor*) applications. It is ignored for all other applications. Time-based One-Time Passwords (TOTPs) are only supported with the authentication method password or public key. For the Service Processor (*service_processor*) application, *none* and *publickey* are the only supported enum values. |

owner

Owner name and UUID that uniquely identifies the user account.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

role

| Name | Type | Description |
| --- | --- | --- |
| name | string | Role name |

account

| Name | Type | Description |
| --- | --- | --- |
| applications | array[account_application] | |
| comment | string | Optional comment for the user account. |
| locked | boolean | Locked status of the account. |
| password | string | Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters. |
| password_hash_algorithm | string | Password hash algorithm used to generate a hash of the user's password for password matching.To modify "password_hash_algorithm", use REST API "/api/security/authentication/password". |
| role | role | |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage the anti-ransomware package

### Manage the anti-ransomware package

#### Overview

This API displays and manages the anti-ransomware package version.

#### Examples

**Retrieving the anti-ransomware package version of all nodes in the cluster.**

The following example shows how to retrieve anti-ransomware package version.

```
# The API:
GET /security/anti-ransomware

# The call:
curl -X GET "https://<mgmt-ip>/api/security/anti-ransomware" -H "accept:
application/hal+json"

# The response:
{
"nodes": [
  {
    "name": "sti11-vsim-ucs573i",
    "version": "1.0"
  },
  {
    "name": "sti11-vsim-ucs573j",
    "version": "1.0"
  }
],
"_links": {
  "self": {
    "href": "/api/security/anti-ransomware/"
  }
}
}
```

### Updating the anti-ransomware package

The following example shows how to update the anti-ransomware package on the cluster.

```
# The API:
PATCH /security/anti-ransomware

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/anti-ransomware" -H "accept:
application/hal+json" -d '{"uri": "http://server/package"}'

# The response:
{
"job": {
  "uuid": "74d0faef-3e1a-11ef-af89-005056ac6d8a",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/74d0faef-3e1a-11ef-af89-005056ac6d8a"
    }
  }
}
}
```

The call to update the anti-ransomware package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the job `state` field is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/74d0faef-3e1a-11ef-af89-
005056ac6d8a" -H "accept: application/hal+json"

# The response:
{
"uuid": "74d0faef-3e1a-11ef-af89-005056ac6d8a",
"description": "PATCH /api/security/anti-ransomware",
"state": "success",
"message": "success",
"code": 0,
"_links": {
  "self": {
    "href": "/api/cluster/jobs/74d0faef-3e1a-11ef-af89-005056ac6d8a"
  }
}
}
```

---

### Retrieve the anti-ransomware version of the nodes in the cluster

GET `/security/anti-ransomware`

**Introduced In:** 9.16

Retrieves the anti-ransomware package version of the nodes in the cluster.

#### Related ONTAP commands

- `security anti-ransomware show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| nodes | array[anti_ransomware_node] | |
| uri | string | URI of the anti-ransomware package through a server |

### Example response

```
{
  "nodes": [
    {
      "name": "node1",
      "version": "1.0"
    }
  ],
  "uri": "http://server/package"
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917775 | Operation is not supported on offline volumes. |
| 203161606 | Failed to check the "anti-ransomware-state" of the volume. |
| 203162108 | Operation is not supported on volumes that contain LUNs or namespaces. |
| 203162109 | Operation is not supported on an anti-ransomware disable-in-progress volume. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

anti_ransomware_node

| Name | Type | Description |
|---|---|---|
| name | string | Name of the node. |
| version | string | Anti-ransomware version. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the anti-ransomware package on the cluster

PATCH `/security/anti-ransomware`

**Introduced In:** 9.16

Updates the anti-ransomware package on the cluster.

#### Required properties

- `uri` - URI of the anti-ransomware package

#### Related ONTAP commands

- `security anti-ransomware update-package-from-uri`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| nodes | array[anti_ransomware_node] | |
| uri | string | URI of the anti-ransomware package through a server |

### Example request

```
{
  "nodes": [
    {
      "name": "node1",
      "version": "1.0"
    }
  ],
  "uri": "http://server/package"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 203161684 | Autonomous Ransomware Protection (ARP) package file must be either a .tar or .tgz file. |
| 203161685 | Autonomous Ransomware Protection (ARP) package file download failed. |
| 203161686 | Autonomous Ransomware Protection (ARP) package file unzip failed. |
| 203161687 | Autonomous Ransomware Protection (ARP) package file untar failed. |
| 203161688 | Autonomous Ransomware Protection (ARP) package archive invalid. No ARP version file found. |
| 203161689 | Internal error. Filesystem error while installing Autonomous Ransomware Protection (ARP) package file. |
| 203161690 | Autonomous Ransomware Protection (ARP) package update requires an effective cluster version of 9.16.1 or later. |
| 203161691 | Autonomous Ransomware Protection (ARP) package file is invalid. The checksum verification of the ARP configuration files failed. |
| 203161692 | Autonomous Ransomware Protection (ARP) package update failed because one or more nodes are not healthy. |
| 203161694 | Failed to update the Autonomous Ransomware Protection (ARP) package file because the `uri` property was not specified. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

anti_ransomware_node

| Name | Type | Description |
|---|---|---|
| name | string | Name of the node. |
| version | string | Anti-ransomware version. |

anti_ransomware

| Name | Type | Description |
|---|---|---|
| nodes | array[anti_ransomware_node] | |
| uri | string | URI of the anti-ransomware package through a server |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage anti-ransomware auto enablement

### Manage anti-ransomware auto enablement

#### Overview

This API provides details about the anti-ransomware auto enablement status including information about warm-up periods and auto settings.

```
# The API
/security/anti-ransomware/auto-enable

# The call:
curl -X GET "https://<mgmt-ip>/api/security/anti-ransomware/auto-enable"
-H "accept: application/hal+json"

# The response:
{
"new_volume_auto_enable": true,
"existing_volume_auto_enable": true,
"warm_up_period_applicable": true,
"warm_up_period_completed": false,
"warm_up_period_duration": "12 hrs",
"warm_up_period_time_remaining" : "6 hrs",
"_links": {
  "self": {
    "href": "/security/anti-ransomware/auto-enable"
  }
}
}
```

#### Retrieve anti-ransomware auto enablement values

GET `/security/anti-ransomware/auto-enable`

**Introduced In:** 9.18

Retrieves the current anti-ransomware auto enablement values, including details about warm-up period and auto enablement setting for new and existing volumes.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| new_volume_auto_enable | boolean | Auto Anti-Ransomware Protection (ARP) setting for new volumes. |
| warm_up_period_applicable | boolean | Indicates if the warm-up period is applicable. |
| warm_up_period_completed | boolean | Indicates if warm-up period has completed. |
| warm_up_period_remaining_duration | string | Time remaining for warm-up period completion. |
| warm_up_period_total_duration | string | Duration of warm-up period. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "warm_up_period_remaining_duration": "string",
  "warm_up_period_total_duration": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Modify the anti-ransomware auto enablement setting

PATCH `/security/anti-ransomware/auto-enable`

**Introduced In:** 9.18

Modifies the anti-ransomware auto enablement setting.

#### Request Body

| Name | Type | Description |
|---|---|---|
| new_volume_auto_enable | boolean | Auto Anti-Ransomware Protection (ARP) setting for new volumes. |

| Name | Type | Description |
| --- | --- | --- |
| warm_up_period_applicable | boolean | Indicates if the warm-up period is applicable. |
| warm_up_period_completed | boolean | Indicates if warm-up period has completed. |
| warm_up_period_remaining_duration | string | Time remaining for warm-up period completion. |
| warm_up_period_total_duration | string | Duration of warm-up period. |

**Example request**

```
{
  "warm_up_period_remaining_duration": "string",
  "warm_up_period_total_duration": "string"
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default, Error
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

anti_ransomware_auto_enable

| Name | Type | Description |
|---|---|---|
| new_volume_auto_enable | boolean | Auto Anti-Ransomware Protection (ARP) setting for new volumes. |
| warm_up_period_applicable | boolean | Indicates if the warm-up period is applicable. |
| warm_up_period_completed | boolean | Indicates if warm-up period has completed. |
| warm_up_period_remaining_duration | string | Time remaining for warm-up period completion. |
| warm_up_period_total_duration | string | Duration of warm-up period. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## View suspect files generated by anti-ransomware

### View suspect files generated by anti-ransomware

#### Retrieving information on suspected files

The suspect GET API retrieves a list of recently suspected files potentially attacked by ransomware.

### Retrieve information on the suspects generated by anti-ransomware analytics

GET `/security/anti-ransomware/suspects`

**Introduced In:** 9.10

Retrieves information on the suspects generated by the anti-ransomware analytics.

#### Related ONTAP commands

- `security anti-ransomware volume attack generate-report`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| file.path | string | query | False | Filter by file.path |
| file.format | string | query | False | Filter by file.format |
| file.suspect_time | string | query | False | Filter by file.suspect_time |
| file.reason | string | query | False | Filter by file.reason<br>• Introduced in: 9.11 |
| file.name | string | query | False | Filter by file.name |
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[anti_ransomware_suspect] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "file": {
        "format": "pdf",
        "name": "test_file",
        "path": "d1/d2/d3",
        "reason": "High Entropy",
        "suspect_time": "2021-05-12 11:00:16 -0400"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

file

| Name | Type | Description |
|---|---|---|
| format | string | File format of the suspected file. |
| name | string | Name of the suspected file. |
| path | string | Path of the suspected file. |
| reason | string | Reason behind this file being suspected |
| suspect_time | string | Time when the file was detected as a potential suspect in date-time format. |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

anti_ransomware_suspect

File suspected to be potentially attacked by ransomware.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| file | file | |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Delete suspect files

### Delete suspect files

The suspects DELETE API clears all the suspect files for a volume from the list of potential suspects.

### Clear suspect files of a volume

DELETE `/security/anti-ransomware/suspects/{volume.uuid}`

**Introduced In:** 9.10

Clears either all the suspect files of a volume or suspect files of a volume based on file format or suspect time provided.

#### Related ONTAP commands

- `security anti-ransomware volume attack clear-suspect`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Identification of the Anti-ransomware suspect file for the deletion.<br>• format: uuid |
| is_false_positive | boolean | query | False | Specifies whether the suspected ransomware activity is a false positive or not.<br>• Default value: |
| file.format | string | query | False | Specifies the file format of the suspect file.<br>• Introduced in: 9.18 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 203161673 | Operation is not supported on a data protection volume. |
| 203162105 | Clearing a suspect by specifying the file format is not supported when an attack is detected by encryption_percentage_analysis. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage data-entropy statistics

## Manage data-entropy statistics

### Overview

This API returns the percentage of data classified as encrypted using an entropy algorithm, measured at various time intervals at the volume level. If no volume UUID is specified, entropy statistics of all the volumes will be returned.
The entropy_stats_type parameter can be used to retrieve statistics with sub_hourly, hourly, or daily granularity. Additionally, it can be used to identify intervals that exhibited high (high_enc_pct) percentages of encrypted
data. If no type is specified, all types of entropy statistics will be returned for the volumes.

### Examples

#### Retrieving volume stats

In this example, the API returns the data-entropy statistics for the volumes.

```
# The API:
/security/anti-ransomware/volume/entropy-stats

# The call:
curl -X GET "https://<mgmt-ip>/api/security/anti-
ransomware/volume/entropy-stats" -H "accept: application/json"

# The response:
{
"records": [
  {
     "volume": {
     "uuid": "61edb8bb-b7d1-11ef-932d-005056bbaeff",
     "name": "v2"
     },
     "entropy_stats_type": "sub_hourly",
     "timestamp": "2024-12-13T03:36:24-05:00",
     "data_written_in_bytes": 13249687,
     "encryption_percentage": 51,
     "duration": "PT10M21S",
     "_links": {
      "self": {
       "href": "/api/security/anti-ransomware/volume/entropy-
stats/61edb8bb-b7d1-11ef-932d-005056bbaeff/sub_hourly/2024-12-
13T03%3A36%3A24-05%3A00"
      }
     }
  },
  {
```

```
      "volume": {
      "uuid": "61edb8bb-b7d1-11ef-932d-005056bbaeff",
      "name": "v2"
      },
      "entropy_stats_type": "hourly",
      "timestamp": "2024-12-13T04:16:06-05:00",
      "data_written_in_bytes": 13249687,
      "encryption_percentage": 51,
      "duration": "PT1H2M36S",
      "_links": {
       "self": {
        "href": "/api/security/anti-ransomware/volume/entropy-
stats/61edb8bb-b7d1-11ef-932d-005056bbaeff/hourly/2024-12-13T04%3A16%3A06-
05%3A00"
       }
      }
  }
],
"num_records": 20,
"_links": {
 "self": {
   "href": "/api/security/anti-ransomware/volume/entropy-stats/61edb8bb-
b7d1-11ef-932d-005056bbaeff"
  }
 }
}
```


**Retrieve data-entropy statistics for volumes**

GET `/security/anti-ransomware/volume/entropy-stats`

**Introduced In:** 9.17

Retrieves the data-entropy statistics for the volumes.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.name | string | query | False | Filter by volume.name |
| volume.uuid | string | query | False | Filter by volume.uuid |
| timestamp | string | query | False | Filter by timestamp |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| data_written_in_bytes | integer | query | False | Filter by data_written_in_bytes |
| encryption_percentage | integer | query | False | Filter by encryption_percentage |
| duration | string | query | False | Filter by duration |
| entropy_stats_type | string | query | False | Filter by entropy_stats_type |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "data_written_in_bytes": 0,
      "duration": "PT15M",
      "encryption_percentage": 0,
      "entropy_stats_type": "string",
      "timestamp": "string",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

records

Volume data-entropy statistics derived from block analysis.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| data_written_in_bytes | integer | The amount of data written. |

| Name | Type | Description |
|---|---|---|
| duration | string | The duration of the interval associated with this statistic. The duration is represented in ISO-8601 standard format. |
| encryption_percentage | integer | The percentage of data that is encrypted. |
| entropy_stats_type | string | Volume metrics are retrieved based on this statistic type. |
| timestamp | string | Start time in date-time format. |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve a data-entropy statistic for a volume**

GET `/security/anti-ransomware/volume/entropy-stats/{volume.uuid}/{entropy_stats_type}/{timestamp}`

**Introduced In:** 9.17

Retrieves a data-entropy statistic for the volumes.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | The UUID of the volume. |
| entropy_stats_type | string | path | True | The type of entropy statistics.<br>• enum: ["sub_hourly", "hourly", "daily", "high_enc_pct"] |
| timestamp | string | path | True | The timestamp of the statistics.<br>• format: date-time |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| data_written_in_bytes | integer | The amount of data written. |
| duration | string | The duration of the interval associated with this statistic. The duration is represented in ISO-8601 standard format. |
| encryption_percentage | integer | The percentage of data that is encrypted. |
| entropy_stats_type | string | Volume metrics are retrieved based on this statistic type. |
| timestamp | string | Start time in date-time format. |
| volume | volume | |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "data_written_in_bytes": 0,
  "duration": "PT15M",
  "encryption_percentage": 0,
  "entropy_stats_type": "string",
  "timestamp": "string",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update audit settings

### View and update audit settings

#### Overview

This API controls what is logged to the audit log files. All operations that make changes are always logged and cannot be disabled. The PATCH request updates administrative audit settings for GET requests. All fields are optional for a PATCH request. A GET request retrieves administrative audit settings for GET requests.
+

#### Examples

### Retrieving administrative audit settings for GET requests

The following example shows the administrative audit settings for GET requests.
+

```
# The API:
/api/security/audit

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit"

# The response:
{
"cli": false,
"http": false,
"ontapi": false,
"_links": {
    "self": {
        "href": "/api/security/audit"
    }
}
}
```

### Updating administrative audit settings for GET requests

The following example updates the administrative audit settings for GET requests
+

```
# The API:
/api/security/audit

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/audit" -d '{"cli": false,
"http": true, "ontapi": true}'
```

### Retrieve administrative audit settings for GET requests

GET `/security/audit`

**Introduced In:** 9.6

Retrieves administrative audit settings for GET requests.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| cli | boolean | Enable auditing of CLI GET Operations. Valid in PATCH |
| http | boolean | Enable auditing of HTTP GET Operations. Valid in PATCH |

| Name | Type | Description |
| --- | --- | --- |
| ontapi | boolean | Enable auditing of ONTAP API GET operations. Valid in PATCH<br>• Introduced in: 9.6<br>• x-nullable: true |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update administrative audit settings for GET requests

PATCH `/security/audit`

**Introduced In:** 9.6

Updates administrative audit settings for GET requests.
All of the fields are optional. An empty body will make no changes.

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| cli | boolean | Enable auditing of CLI GET Operations. Valid in PATCH |
| http | boolean | Enable auditing of HTTP GET Operations. Valid in PATCH |
| ontapi | boolean | Enable auditing of ONTAP API GET operations. Valid in PATCH<br>• Introduced in: 9.6<br>• x-nullable: true |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| cli | boolean | Enable auditing of CLI GET Operations. Valid in PATCH |
| http | boolean | Enable auditing of HTTP GET Operations. Valid in PATCH |
| ontapi | boolean | Enable auditing of ONTAP API GET operations. Valid in PATCH<br>• Introduced in: 9.6<br>• x-nullable: true |

**Error**

```
Status: Default, Error
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

security_audit

| Name | Type | Description |
|---|---|---|
| cli | boolean | Enable auditing of CLI GET Operations. Valid in PATCH |
| http | boolean | Enable auditing of HTTP GET Operations. Valid in PATCH |
| ontapi | boolean | Enable auditing of ONTAP API GET operations. Valid in PATCH<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Forward audit logs to syslog/splunk servers

## Forward audit logs to syslog/splunk servers

### Overview

This API controls the forwarding of audit log information to remote syslog/splunk servers. Multiple destinations can be configured and all audit records are forwarded to all destinations.

A GET operation retrieves information about remote syslog/splunk server destinations.
A POST operation creates a remote syslog/splunk server destination.
A GET operation on /security/audit/destinations/{address}/{port} retrieves information about the syslog/splunk server destination given its address and port number.
A PATCH operation on /security/audit/destinations/{address}/{port} updates information about the syslog/splunk server destination given its address and port number.
A DELETE operation on /security/audit/destinations/{address}/{port} deletes a syslog/splunk server destination given its address and port number.

## Overview of fields used for creating a remote syslog/splunk destination

The fields used for creating a remote syslog/splunk destination fall into the following categories

## Required properties

All of the following fields are required for creating a remote syslog/splunk destination

- `address`

## Optional properties

All of the following fields are optional for creating a remote syslog/splunk destination

- `port`
- `ipspace`
- `protocol`
- `facility`
- `verify_server`
- `message_format` (Can be either "legacy_netapp" or "rfc_5424")
- `timestamp_format_override` (Can be either "no_override", "rfc_3164", "iso_8601_utc" or "iso_8601_local_time")
- `hostname_format_override` (Can be either "no_override", "fqdn" or "hostname_only")
  +

---

### Examples

## Retrieving remote syslog/splunk server destinations

The following example shows remote syslog/splunk server destinations

+

```
# The API:
/api/security/audit/destinations

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit/destinations"

# The response:
{
"records": [
  {
    "address": "1.1.1.1",
    "port": 514,
    "_links": {
      "self": {
        "href": "/api/security/audit/destinations/1.1.1.1/514"
      }
  }
}
 ],
 "num_records": 1,
 "_links": {
   "self": {
     "href": "/api/security/audit/destinations"
   }
 }
}
```

#### Creating remote syslog/splunk server destinations

The following example creates remote syslog/splunk server destinations.
+

```
# The API:
/api/security/audit/destinations

# The call:
curl -X POST "https://<cluster-
ip>/api/security/audit/destinations?force=true" -d '{ "address":
"1.1.1.1", "port": 514, "protocol": "udp_unencrypted", "facility":
"kern"}'
```

---

#### Retrieving a remote syslog/splunk server destination given its destination address and port number

The following example retrieves a remote syslog/splunk server destination given its destination address and port number.
+

---

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X GET "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514"

# The response:
{
"address": "1.1.1.1",
"port": 514,
"ipspace": {
  "name": "Default",
  "uuid": "a97a3549-f7ae-11ec-b6bc-005056a7c8ff"
},
"protocol": "udp_unencrypted",
"facility": "kern",
"verify_server": false,
"message_format": "legacy_netapp",
"timestamp_format_override": "no_override",
"hostname_format_override": "no_override",
"_links": {
  "self": {
    "href": "/api/security/audit/destinations/1.1.1.1/514"
  }
}
}
```

---

#### Updating a remote syslog/splunk server destination given its destination address and port number

The following example updates a remote syslog/splunk server destination configuration given its destination address and port number.
+

---

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X PATCH "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514" -d '{"facility":
"user"}'
```

#### Deleting a remote syslog/splunk server destination given its destination address and port number

The following example deletes a remote syslog/splunk server destination configuration given its destination address and port number.
+

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X DELETE "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514"
```

### Define a remote syslog or splunk server to receive audit information

GET `/security/audit/destinations`

**Introduced In:** 9.6

Defines a remote syslog/splunk server for sending audit information to.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| port | integer | query | False | Filter by port |
| message_format | string | query | False | Filter by message_format<br>• Introduced in: 9.13 |
| timestamp_format_override | string | query | False | Filter by timestamp_format_override<br>• Introduced in: 9.13 |
| protocol | string | query | False | Filter by protocol |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| ipspace.name | string | query | False | Filter by ipspace.name<br>• Introduced in: 9.12 |
| ipspace.uuid | string | query | False | Filter by ipspace.uuid<br>• Introduced in: 9.12 |
| address | string | query | False | Filter by address |
| facility | string | query | False | Filter by facility |
| verify_server | boolean | query | False | Filter by verify_server |
| hostname_format_override | string | query | False | Filter by hostname_format_override<br>• Introduced in: 9.13 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |
| Limit the number of records returned. | return_timeout | integer | query | False |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[security_audit_log_forward] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "address": "1.1.1.1",
      "facility": "string",
      "hostname_format_override": "string",
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "message_format": "string",
      "port": 514,
      "protocol": "string",
      "timestamp_format_override": "string"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

security_audit_log_forward

| Name | Type | Description |
|---|---|---|
| address | string | Destination syslog\|splunk host to forward audit records to. This can be an IP address (IPv4\|IPv6) or a hostname. |
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |

| Name | Type | Description |
| --- | --- | --- |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| port | integer | Destination Port. The default port depends on the protocol chosen:<br>For un-encrypted destinations the default port is 514.<br>For encrypted destinations the default port is 6514. |
| protocol | string | Log forwarding protocol |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

**Define the remote syslog or splunk server information**

POST `/security/audit/destinations`

**Introduced In:** 9.6

Configures remote syslog/splunk server information.

#### Required properties

All of the following fields are required for creating a remote syslog/splunk destination

- `address`

#### Optional properties

All of the following fields are optional for creating a remote syslog/splunk destination

- `port` (1 - 65535)
- `ipspace`
- `protocol`
- `facility`
- `verify_server` (Can only be "true" when protocol is "tcp_encrypted")
- `message_format` (Can be either "legacy-netapp" or "rfc-5424")
- `timestamp_format_override` (Can be either "no-override", "rfc-3164", "iso-8601-utc" or "iso-8601-local-time")
- `hostname_format_override` (Can be either "no-override", "fqdn" or "hostname-only")

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| force | boolean | query | False | Skip the Connectivity Test<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| address | string | Destination syslog\|splunk host to forward audit records to. This can be an IP address (IPv4\|IPv6) or a hostname. |

| Name | Type | Description |
| --- | --- | --- |
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| port | integer | Destination Port. The default port depends on the protocol chosen:<br>For un-encrypted destinations the default port is 514.<br>For encrypted destinations the default port is 6514. |
| protocol | string | Log forwarding protocol |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server’s certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server’s certificate. Setting "verify_server" to "false" will not enforce validation of remote server’s certificate. |

### Example request

```
{
  "address": "1.1.1.1",
  "facility": "string",
  "hostname_format_override": "string",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "message_format": "string",
  "port": 514,
  "protocol": "string",
  "timestamp_format_override": "string"
}
```

### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[security_audit_log_forward] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "address": "1.1.1.1",
      "facility": "string",
      "hostname_format_override": "string",
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "message_format": "string",
      "port": 514,
      "protocol": "string",
      "timestamp_format_override": "string"
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 15661 | The object specified could not be found |
| 13114 | Internal error |
| 13115 | Invalid input |

| Error Code | Description |
| --- | --- |
| 4522285 | Server verification cannot be enabled because it requires a protocol with encryption. Encryption can be selected using the protocol field. |
| 9240603 | Cannot ping destination host. Verify connectivity to desired host or skip the connectivity check with the -force parameter. |
| 327698 | Failed to create RPC client to destination host |
| 9240609 | Cannot connect to destination host. |
| 9240604 | Cannot resolve the destination host. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

security_audit_log_forward

| Name | Type | Description |
|---|---|---|
| address | string | Destination syslog\|splunk host to forward audit records to. This can be an IP address (IPv4\|IPv6) or a hostname. |
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| port | integer | Destination Port. The default port depends on the protocol chosen:<br>For un-encrypted destinations the default port is 514.<br>For encrypted destinations the default port is 6514. |
| protocol | string | Log forwarding protocol |

| Name | Type | Description |
|---|---|---|
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Delete the remote syslog or splunk server information**

DELETE `/security/audit/destinations/{address}/{port}`

**Introduced In:** 9.6

Deletes remote syslog/splunk server information.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| address | string | path | True | IP address of remote syslog/splunk server. |
| port | integer | path | True | Port number of remote syslog/splunk server. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the remote syslog or splunk server information

GET `/security/audit/destinations/{address}/{port}`

**Introduced In:** 9.6

Defines a remote syslog/splunk server for sending audit information to.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| address | string | path | True | IP address of remote syslog/splunk server. |
| port | integer | path | True | Port number of remote syslog/splunk server. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| address | string | Destination syslog\|splunk host to forward audit records to. This can be an IP address (IPv4\|IPv6) or a hostname. |
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| port | integer | Destination Port. The default port depends on the protocol chosen:<br>For un-encrypted destinations the default port is 514.<br>For encrypted destinations the default port is 6514. |
| protocol | string | Log forwarding protocol |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

#### Example response

```
{
  "address": "1.1.1.1",
  "facility": "string",
  "hostname_format_override": "string",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "message_format": "string",
  "port": 514,
  "protocol": "string",
  "timestamp_format_override": "string"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipspace

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | IPspace name |
| uuid | string | IPspace UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the remote syslog or splunk server information

PATCH `/security/audit/destinations/{address}/{port}`

**Introduced In:** 9.6

Updates remote syslog/splunk server information.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| address | string | path | True | IP address of remote syslog/splunk server. |
| port | integer | path | True | Port number of remote syslog/splunk server. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

#### Example request

```
{
  "facility": "string",
  "hostname_format_override": "string",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "message_format": "string",
  "timestamp_format_override": "string"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

### Example response

```
{
  "facility": "string",
  "hostname_format_override": "string",
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "message_format": "string",
  "timestamp_format_override": "string"
}
```

### Error

```
Status: Default, Default
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipspace

| Name | Type | Description |
|---|---|---|
| name | string | IPspace name |
| uuid | string | IPspace UUID |

security_audit_log_forward

| Name | Type | Description |
|---|---|---|
| facility | string | This is the standard Syslog Facility value that is used when sending audit records to a remote server. |
| hostname_format_override | string | Syslog Hostname Format Override |
| ipspace | ipspace | |
| message_format | string | Syslog message format to be used. legacy_netapp format (variation of RFC-3164) is default message format. |
| timestamp_format_override | string | Syslog Timestamp Format Override. |
| verify_server | boolean | This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View administrative audit logs

### View administrative audit logs

#### Overview

These APIs return audit log records. The GET requests retrieves all audit log records. An audit log record contains information such as timestamp, node name, index and so on.
+

---

#### Example

### Retrieving audit log records

The following example shows the audit log records.
Note: The index field is used to order the audit log messages before they are displayed. If multiple entries for the same node and timestamp occur simultaneously, the index assigns an order to ensure logical consistency.
+

---

```
# The API:
/api/security/audit/messages

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit/messages"

# The response:
{
"records": [
  {
    "timestamp": "2019-03-08T11:03:32-05:00",
    "node": {
      "name": "node1",
      "uuid": "bc9af9da-41bb-11e9-a3db-005056bb27cf",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/bc9af9da-41bb-11e9-a3db-
005056bb27cf"
        }
      }
    },
    "index": 4294967299,
    "application": "http",
    "location": "172.21.16.89",
    "user": "admin",
    "input": "GET /api/security/audit/destinations/",
    "state": "pending",
    "scope": "cluster",
    "role": "admin"
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/audit/messages"
  }
}
}
```

---

**Retrieve the administrative audit log viewer**

GET `/security/audit/messages`

**Introduced In:** 9.6

---

Retrieves the administrative audit log viewer.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| user | string | query | False | Filter by user |
| location | string | query | False | Filter by location |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| state | string | query | False | Filter by state |
| command_id | string | query | False | Filter by command_id |
| timestamp | string | query | False | Filter by timestamp |
| application | string | query | False | Filter by application |
| scope | string | query | False | Filter by scope |
| index | integer | query | False | Filter by index |
| message | string | query | False | Filter by message |
| session_id | string | query | False | Filter by session_id |
| input | string | query | False | Filter by input |
| svm.name | string | query | False | Filter by svm.name |
| role | string | query | False | Filter by role<br>• Introduced in: 9.17 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[security_audit_log] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": "string",
      "command_id": "string",
      "index": 0,
      "input": "string",
      "location": "string",
      "message": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "role": "string",
      "scope": "string",
      "session_id": "string",
      "state": "string",
      "svm": {
        "name": "string"
      },
      "timestamp": "string",
      "user": "string"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

Node where the audit message resides.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

svm

This is the SVM through which the user connected.

| Name | Type | Description |
|---|---|---|
| name | string | |

security_audit_log

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| application | string | This identifies the "application" by which the request was processed. |

| Name | Type | Description |
| --- | --- | --- |
| command_id | string | This is the command ID for this request.<br>Each command received on a CLI session is assigned a command ID. This enables you to correlate a request and response. |
| index | integer | Internal index for accessing records with the same time and node. This is a 64-bit unsigned value that is used to order the audit log messages before they are displayed. If multiple entries for the same node and timestamp occur simultaneously, the index assigns an order to ensure logical consistency. |
| input | string | The request. |
| location | string | This identifies the location of the remote user. This is an IP address or "console". |
| message | string | This is an optional field that might contain "error" or "additional information" about the status of a command. |
| node | node | Node where the audit message resides. |
| role | string | Role of the remote user. |
| scope | string | Set to "svm" when the request is on a data SVM; otherwise set to "cluster". |
| session_id | string | This is the session ID on which the request is received. Each SSH session is assigned a session ID.<br>Each http/ontapi/snmp request is assigned a unique session ID. |
| state | string | State of of this request. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | This is the SVM through which the user connected. |
| timestamp | string | Log entry timestamp. Valid in URL |
| user | string | Username of the remote user. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage data SVM account information

### Manage data SVM account information

#### Overview

This API configures data SVM account information at the Active Directory. For Active Directory domain-based authentication for cluster accounts, a data SVM must be configured and registered as a machine account at the Active Directory. All authentication requests are proxied through this SVM.

#### Examples

### Creating a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
POST  "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ad-
proxy" -d '{"svm.uuid":"13f87d78-70c7-11e9-b722-0050568ec89f"}'
```

#### Updating a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
PATCH "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/ad-
proxy" -d '{"svm.uuid":"13f87d78-70c7-11e9-b722-0050568ec89f"}'
```

#### Retrieving a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
GET "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/ad-
proxy"

# The response:
{
"svm": {
  "uuid": "512eab7a-6bf9-11e9-a896-005056bb9ce1",
  "name": "vs2",
  "_links": {
    "self": {
      "href": "/api/svm/svms/512eab7a-6bf9-11e9-a896-005056bb9ce1"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/security/authentication/cluster/ad-proxy"
  }
}
}
```

### Delete a data SVM configured as a tunnel

DELETE `/security/authentication/cluster/ad-proxy`

**Introduced In:** 9.7

Deletes the data SVM configured as a tunnel for Active Directory based authentication for cluster user accounts.

#### Related ONTAP commands

- `security login domain-tunnel delete`

#### Learn more

- DOC /security/authentication/cluster/ad-proxy
- DOC /security/accounts

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve SVM information configured as an Active Directory domain-tunnel

GET `/security/authentication/cluster/ad-proxy`

**Introduced In:** 9.7

Retrieves SVM information configured as an Active Directory domain-tunnel.

#### Related ONTAP commands

- `security login domain-tunnel show`

#### Learn more

- DOC /security/authentication/cluster/ad-proxy
- DOC /security/accounts

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a data SVM configured as a tunnel

PATCH `/security/authentication/cluster/ad-proxy`

**Introduced In:** 9.7

Updates the data SVM configured as a tunnel for Active Directory based authentication for cluster user accounts.

#### Related ONTAP commands

- `security login domain-tunnel modify`

#### Learn more

- DOC /security/authentication/cluster/ad-proxy
- DOC /security/accounts

#### Request Body

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

#### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 2621516 | This operation is only supported on a data SVM. |
| 2621706 | The specified `svm.uuid` and `svm.name` refer to different SVMs. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cluster_ad_proxy

The SVM configured as proxy for Active Directory authentication of cluster accounts.

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Configure a data SVM as a proxy

POST `/security/authentication/cluster/ad-proxy`

**Introduced In:** 9.7

Configures a data SVM as a proxy for Active Directory based authentication for cluster user accounts.

#### Required properties

- `svm.name` or `svm.uuid` - Name and UUID of the SVM for a cluster user account.

#### Related ONTAP commands

- `security login domain-tunnel create`

#### Learn more

- DOC /security/authentication/cluster/ad-proxy
- DOC /security/accounts

#### Request Body

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

### Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621516 | This operation is only supported on a data SVM. |
| 2621706 | The specified `svm.uuid` and `svm.name` refer to different SVMs. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

cluster_ad_proxy

The SVM configured as proxy for Active Directory authentication of cluster accounts.

| Name | Type | Description |
|---|---|---|
| svm | svm | SVM, applies only to SVM-scoped objects. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage LDAP server configuration

### Manage LDAP server configuration

#### Overview

LDAP servers are used to centrally maintain user information. LDAP configurations must be set up to look up information stored in the LDAP directory on the external LDAP servers. This API is used to retrieve and manage
cluster LDAP server configurations.

#### Examples

### Retrieving the cluster LDAP information

The cluster LDAP GET request retrieves the LDAP configuration of the cluster.

The following example shows how a GET request is used to retrieve the cluster LDAP information:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json"

# The response:
{
"servers": [
  "10.10.10.10",
  "domainB.example.com"
],
"schema": "ad_idmu",
"port": 389,
"min_bind_level": "anonymous",
"bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
"base_dn": "dc=domainA,dc=example,dc=com",
"base_scope": "subtree",
"use_start_tls": true,
"session_security": "none",
"try_channel_binding": true,
"_links": {
  "self": {
    "href": "/api/security/authentication/cluster/ldap"
  }
}
}
```

**Creating the cluster LDAP configuration**

The cluster LDAP POST operation creates an LDAP configuration for the cluster.

The following example shows how to issue a POST request with all of the fields specified:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json" -H "Content-Type: application/json" -d
"{ \"servers\": [ \"10.10.10.10\", \"domainB.example.com\" ], \"schema\":
\"ad_idmu\", \"port\": 389, \"min_bind_level\": \"anonymous\",
\"bind_dn\": \"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"<BIND-PASSWORD>\", \"base_dn\":
\"dc=domainA,dc=example,dc=com\", \"base_scope\": \"subtree\",
\"use_start_tls\": false, \"session_security\": \"none\"}"
```

The following example shows how to issue a POST request with a number of optional fields not specified:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json" -H "Content-Type: application/json" -d
"{ \"port\": 389, \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"<BIND-PASSWORD>\", \"base_dn\":
\"dc=domainA,dc=example,dc=com\", \"session_security\": \"none\"}"
```

#### Updating the cluster LDAP configuration

The cluster LDAP PATCH request updates the LDAP configuration of the cluster.

The following example shows how a PATCH request is used to update the cluster LDAP configuration:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"servers\": [ \"55.55.55.55\" ], \"schema\": \"ad_idmu\", \"port\": 636,
\"use_start_tls\": false }"
```

#### Deleting the cluster LDAP configuration

The cluster LDAP DELETE request deletes the LDAP configuration of the cluster.

The following example shows how a DELETE request is used to delete the cluster LDAP configuration:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/security/authentication/cluster/ldap" -H "accept:
application/hal+json"
```

**Delete the LDAP configuration for the cluster**

DELETE `/security/authentication/cluster/ldap`

**Introduced In:** 9.6

Deletes the LDAP configuration of the cluster.

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the LDAP configuration for the cluster

GET `/security/authentication/cluster/ldap`

**Introduced In:** 9.6

Retrieves the cluster LDAP configuration.

#### Related ONTAP commands

- `ldap show`
- `ldap check -vserver vs0`
- `ldap check-ipv6 -vserver vs0`

#### Important notes

- The status.code, status.dn_message, status.message, and status.state fields have the same status fields that are returned using the "ldap check" CLI command.
- Refer to the ipv4 or ipv6 objects available in the status field to get specific information about the code, dn_messages, or message and state information for ipv4 or ipv6.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

| Name | Type | Description |
| --- | --- | --- |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "group_dn": "string",
  "group_membership_filter": "string",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "string",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "string",
  "netgroup_scope": "string",
  "port": 389,
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "skip_config_validation": true,
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv4_state": "string",
    "ipv6": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
```

```
        "state": "string"
      },
      "ipv6_state": "string",
      "message": "string",
      "state": "string"
    },
    "user_dn": "string",
    "user_scope": "string"
  }
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |

| Name | Type | Description |
| --- | --- | --- |
| dn_message | array[string] | |
| ipv4 | ipv4 | |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the LDAP configuration for the cluster

PATCH `/security/authentication/cluster/ldap`

**Introduced In:** 9.6

Both mandatory and optional parameters of the LDAP configuration can be updated.
IPv6 must be enabled if IPv6 family addresses are specified. Configuring more than one LDAP server is recommended to avoid a single point of failure. Both FQDNs and IP addresses are supported for the `servers` property.
The LDAP servers are validated as part of this operation. LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server is invalid.
3. The server is unreachable.

#### Request Body

| Name | Type | Description |
|---|---|---|
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |

| Name | Type | Description |
| --- | --- | --- |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
| --- | --- | --- |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example request**

```
{
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "group_dn": "string",
  "group_membership_filter": "string",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "string",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "string",
  "netgroup_scope": "string",
  "port": 389,
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv4_state": "string",
    "ipv6": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv6_state": "string",
    "message": "string",
    "state": "string"
```

```
  },
  "user_dn": "string",
  "user_scope": "string"
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4915203 | The specified LDAP schema does not exist. |
| 262222 | The specified LDAP servers contain duplicate server entries. |
| 4915229 | DNS resolution failed due to an internal error. Contact technical support if this issue persists. |
| 4915231 | DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured. |
| 23724132 | DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured. |
| 4915234 | Specified LDAP server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast. |
| 4915248 | LDAP servers cannot be empty or "-". Specified FQDN is not valid because it is empty or "-" or it contains either special characters or "-" at the start or end of the domain. |
| 4915251 | STARTTLS and LDAPS cannot be used together |
| 4915257 | The LDAP configuration is not valid. Verify that the Distinguished Names and bind password are correct. |
| 4915258 | The LDAP configuration is not valid. Verify that the servers are reachable and that the network configuration is correct. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces. |

| Error Code | Description |
| --- | --- |
| 4915252 | LDAP referral is not supported with STARTTLS, with session security levels sign, seal or with LDAPS. |
| 4915244 | RPC failure occurred during validation of the LDAP configuration. |
| 4915265 | The specified bind password or bind DN is invalid |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |
| dn_message | array[string] | |
| ipv4 | ipv4 | |

| Name | Type | Description |
| --- | --- | --- |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

cluster_ldap

| Name | Type | Description |
| --- | --- | --- |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |

| Name | Type | Description |
| --- | --- | --- |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
|---|---|---|
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create the LDAP configuration for the cluster**

POST `/security/authentication/cluster/ldap`

**Introduced In:** 9.6

A cluster can have only one LDAP configuration. IPv6 must be enabled if IPv6 family addresses are specified.

**Required properties**

- `servers` - List of LDAP servers used for this client configuration.

#### Recommended optional properties

- `schema` - Schema template name.
- `port` - Port used to connect to the LDAP Servers.
- `ldaps_enabled` - Specifies whether or not LDAPS is enabled.
- `min_bind_level` - Minimum bind authentication level.
- `bind_dn` - Specifies the user that binds to the LDAP servers.
- `base_dn` - Specifies the default base DN for all searches.
- `bind_password` - Specifies the bind password for the LDAP servers.
- `base_scope` - Specifies the default search scope for LDAP queries.
- `use_start_tls` - Specifies whether or not to use Start TLS over LDAP connections.
- `session_security` - Specifies the level of security to be used for LDAP communications.
- `bind_as_cifs_server` - Indicates if CIFS server's credentials are used to bind to the LDAP server.
- `query_timeout` - Maximum time to wait for a query response from the LDAP server, in seconds.
- `user_dn` - User Distinguished Name (DN) used as the starting point in the LDAP directory tree for user lookups.
- `user_scope` - Default search scope for LDAP for user lookups.
- `group_dn` - Group Distinguished Name (DN) used as the starting point in the LDAP directory tree for group lookups.
- `group_scope` - Default search scope for LDAP for group lookups.
- `netgroup_dn` - Netgroup Distinguished Name (DN) used as the starting point in the LDAP directory tree for netgroup lookups.
- `netgroup_scope` - Default search scope for LDAP for netgroup lookups.
- `netgroup_byhost_dn` - Netgroup Distinguished Name (DN) used as the starting point in the LDAP directory tree for netgroup by host lookups.
- `netgroup_byhost_scope` - Default search scope for LDAP for netgroup by host lookups.
- `is_netgroup_byhost_enabled` - Specifies whether netgroup by host querying is enabled.
- `group_membership_filter` - Custom filter used for group membership lookup from an LDAP server.
- `skip_config_validation` - Indicates whether or not the validation for the specified LDAP configuration is disabled.

#### Default property values

- `schema` - *RFC-2307*
- `port` - *389*
- `ldaps_enabled` - *false*
- `min_bind_level` - *simple*
- `base_scope` - *subtree*

- `use_start_tls` - *false*
- `session_security` - *none*
- `query_timeout` - *3*
- `user_scope` - *subtree*
- `group_scope` - *subtree*
- `netgroup_scope` - *subtree*
- `netgroup_byhost_scope` - *subtree*
- `is_netgroup_byhost_enabled` - *false*
- `skip_config_validation` - *false*
- `try_channel_binding` - *true*

Configuring more than one LDAP server is recommended to avoid a single point of failure. Both FQDNs and IP addresses are supported for the `servers` property.
The LDAP servers are validated as part of this operation. LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server is invalid.
3. The server is unreachable.

**Request Body**

| Name | Type | Description |
|---|---|---|
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |

| Name | Type | Description |
| --- | --- | --- |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
| --- | --- | --- |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

**Example request**

```
{
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "group_dn": "string",
  "group_membership_filter": "string",
  "group_scope": "string",
  "min_bind_level": "string",
  "netgroup_byhost_dn": "string",
  "netgroup_byhost_scope": "string",
  "netgroup_dn": "string",
  "netgroup_scope": "string",
  "port": 389,
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "status": {
    "code": 65537300,
    "dn_message": [
      "string"
    ],
    "ipv4": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv4_state": "string",
    "ipv6": {
      "code": 65537300,
      "dn_messages": [
        "string"
      ],
      "message": "string",
      "state": "string"
    },
    "ipv6_state": "string",
    "message": "string",
    "state": "string"
```

```
    },
    "user_dn": "string",
    "user_scope": "string"
  }
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of LDAP records. |
| records | array[ldap_service] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "ad_domain": "example.com",
      "base_dn": "dc=domainB,dc=example,dc=com",
      "base_scope": "string",
      "bind_dn":
"cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
      "bind_password": "abc",
      "group_dn": "cn=abc,users,dc=com",
      "group_membership_filter": "",
      "group_scope": "string",
      "min_bind_level": "string",
      "netgroup_byhost_dn": "cn=abc,users,dc=com",
      "netgroup_byhost_scope": "string",
      "netgroup_dn": "cn=abc,users,dc=com",
      "netgroup_scope": "string",
      "port": 389,
      "preferred_ad_servers": [
        "11.11.11.11"
      ],
      "schema": "ad_idmu",
      "servers": [
        [
          "10.10.10.10",
          "domainB.example.com"
        ]
      ],
      "session_security": "string",
      "status": {
        "code": 65537300,
        "dn_message": [
          "string"
        ],
        "ipv4": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv4_state": "string",
```

```
        "ipv6": {
          "code": 65537300,
          "dn_messages": [
            "string"
          ],
          "message": "string",
          "state": "string"
        },
        "ipv6_state": "string",
        "message": "string",
        "state": "string"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "user_dn": "cn=abc,users,dc=com",
      "user_scope": "string"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 4915203 | The specified LDAP schema does not exist. |
| 262222 | The specified LDAP servers contain duplicate server entries. |
| 4915229 | DNS resolution failed due to an internal error. Contact technical support if this issue persists. |
| 4915231 | DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured. |

| Error Code | Description |
| --- | --- |
| 23724132 | DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured. |
| 4915234 | The specified LDAP server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast. |
| 4915248 | LDAP servers cannot be empty or "-". Specified FQDN is invalid because it is empty or "-" or it contains either special characters or "-" at the start or end of the domain. |
| 4915251 | STARTTLS and LDAPS cannot be used together. |
| 4915257 | The LDAP configuration is invalid. Verify that bind-dn and bind password are correct. |
| 4915258 | The LDAP configuration is invalid. Verify that the servers are reachable and that the network configuration is correct. |
| 13434916 | The SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces. |
| 4915252 | LDAP referral is not supported with STARTTLS, with session security levels sign, seal or with LDAPS. |
| 4915265 | The specified bind password or bind DN is invalid |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ipv4

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

ipv6

| Name | Type | Description |
|---|---|---|
| code | integer | Code corresponding to the error message. If there is no error, it is 0 to indicate success. |
| dn_messages | array[string] | |
| message | string | Provides additional details on the error. |
| state | string | Status of the LDAP service. |

status

| Name | Type | Description |
|---|---|---|
| code | integer | This field is no longer supported. Use ipv4.code or ipv6.code instead. |
| dn_message | array[string] | |
| ipv4 | ipv4 | |

| Name | Type | Description |
| --- | --- | --- |
| ipv4_state | string | This field is no longer supported. Use ipv4.state instead. |
| ipv6 | ipv6 | |
| ipv6_state | string | This field is no longer supported. Use ipv6.state instead. |
| message | string | This field is no longer supported. Use ipv4.message or ipv6.message instead. |
| state | string | The status of the LDAP service for the SVM. The LDAP service is up if either `ipv4_state` or `ipv6_state` is up.<br>The LDAP service is down if both `ipv4_state` and `ipv6_state` are down. |

cluster_ldap

| Name | Type | Description |
| --- | --- | --- |
| base_dn | string | Specifies the default base DN for all searches. |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |

| Name | Type | Description |
| --- | --- | --- |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |

| Name | Type | Description |
| --- | --- | --- |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |

| Name | Type | Description |
| --- | --- | --- |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

ldap_service

| Name | Type | Description |
| --- | --- | --- |
| ad_domain | string | This parameter specifies the name of the Active Directory domain<br>used to discover LDAP servers for use by this client.<br>This is mutually exclusive with `servers` during POST and PATCH. |
| base_dn | string | Specifies the default base DN for all searches. |

| Name | Type | Description |
| --- | --- | --- |
| base_scope | string | Specifies the default search scope for LDAP queries:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| bind_as_cifs_server | boolean | Specifies whether or not CIFS server's credentials are used to bind to the LDAP server. |
| bind_dn | string | Specifies the user that binds to the LDAP servers. |
| bind_password | string | Specifies the bind password for the LDAP servers. |
| group_dn | string | Specifies the group Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for group lookups. |
| group_membership_filter | string | Specifies the custom filter used for group membership lookups from an LDAP server. |
| group_scope | string | Specifies the default search scope for LDAP for group lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| is_netgroup_byhost_enabled | boolean | Specifies whether or not netgroup by host querying is enabled. |

| Name | Type | Description |
| --- | --- | --- |
| is_owner | boolean | Specifies whether or not the SVM owns the LDAP client configuration. |
| ldaps_enabled | boolean | Specifies whether or not LDAPS is enabled. |
| min_bind_level | string | The minimum bind authentication level. Possible values are:<br>• anonymous - anonymous bind<br>• simple - simple bind<br>• sasl - Simple Authentication and Security Layer (SASL) bind |
| netgroup_byhost_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup by host lookups. |
| netgroup_byhost_scope | string | Specifies the default search scope for LDAP for netgroup by host lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| netgroup_dn | string | Specifies the netgroup Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for netgroup lookups. |

| Name | Type | Description |
| --- | --- | --- |
| netgroup_scope | string | Specifies the default search scope for LDAP for netgroup lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |
| port | integer | The port used to connect to the LDAP Servers. |
| preferred_ad_servers | array[string] | |
| query_timeout | integer | Specifies the maximum time to wait for a query response from the LDAP server, in seconds. |
| referral_enabled | boolean | Specifies whether or not LDAP referral is enabled. |
| restrict_discovery_to_site | boolean | Specifies whether or not LDAP server discovery is restricted to site-scope. |
| schema | string | The name of the schema template used by the SVM.<br>• AD-IDMU - Active Directory Identity Management for UNIX<br>• AD-SFU - Active Directory Services for UNIX<br>• MS-AD-BIS - Active Directory Identity Management for UNIX<br>• RFC-2307 - Schema based on RFC 2307<br>• Custom schema |
| servers | array[string] | |

| Name | Type | Description |
| --- | --- | --- |
| session_security | string | Specifies the level of security to be used for LDAP communications:<br>• none - no signing or sealing<br>• sign - sign LDAP traffic<br>• seal - seal and sign LDAP traffic |
| skip_config_validation | boolean | Indicates whether or not the validation for the specified LDAP configuration is disabled. |
| status | status | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| try_channel_binding | boolean | Specifies whether or not channel binding is attempted in the case of TLS/LDAPS. |
| use_start_tls | boolean | Specifies whether or not to use Start TLS over LDAP connections. |
| user_dn | string | Specifies the user Distinguished Name (DN) that is used as the starting point in the LDAP directory tree for user lookups. |
| user_scope | string | Specifies the default search scope for LDAP for user lookups:<br>• base - search the named entry only<br>• onelevel - search all entries immediately below the DN<br>• subtree - search the named DN entry and the entire subtree below the DN |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage NIS configuration

### Manage NIS configuration

#### Overview

NIS servers are used to authenticate user and client computers. NIS domain name and NIS server information is required to configure NIS. This API retrieves and manages NIS server configurations.

#### Examples

### Retrieving cluster NIS information

The cluster NIS GET request retrieves the NIS configuration of the cluster.

The following example shows how a GET request is used to retrieve the cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/nis" -H
"accept: application/hal+json"

# The response:
{
"domain": "domainA.example.com",
"servers": [
  "10.10.10.10",
  "example.com"
],
"bound_servers": [
  "10.10.10.10"
]
}
```

#### Creating the cluster NIS configuration

The cluster NIS POST request creates a NIS configuration for the cluster.

The following example shows how a POST request is used to create a cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"domain\": \"domainA.example.com\", \"servers\": [
\"10.10.10.10\",\"example.com\" ]}"
```

#### Updating the cluster NIS configuration

The cluster NIS PATCH request updates the NIS configuration of the cluster.

The following example shows how to update the domain:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"domain\": \"domainC.example.com\", \"servers\": [ \"13.13.13.13\" ]}"
```

The following example shows how to update the server:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"servers\": [ \"14.14.14.14\" ]}"
```

#### Deleting the cluster NIS configuration

The cluster NIS DELETE request deletes the NIS configuration of the cluster.

The following example shows how a DELETE request is used to delete the cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/hal+json"
```

---

### Delete the NIS configuration for the cluster

DELETE `/security/authentication/cluster/nis`

**Introduced In:** 9.6

Deletes the NIS configuration of the cluster. NIS can be removed as a source from ns-switch if NIS is not used for lookups.

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the NIS configuration for the cluster

GET `/security/authentication/cluster/nis`

**Introduced In:** 9.6

Retrieves the NIS configuration of the cluster. Both NIS domain and servers are displayed by default. The `bound_servers` property indicates the successfully bound NIS servers.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
|---|---|---|
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the NIS configuration for the cluster

PATCH `/security/authentication/cluster/nis`

**Introduced In:** 9.6

Both NIS domain and servers can be updated. Domains and servers cannot be empty. Both FQDNs and IP addresses are supported for the 'servers' field. If the domain is updated, NIS servers must also be specified. IPv6 must be enabled if IPv6 family addresses are specified for the `servers` property.

#### Request Body

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

#### Example request

```
{
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966253 | IPv6 is not enabled in the cluster . |
| 3276964 | The NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters. |
| 3276933 | A maximum of 10 NIS servers can be configured per SVM. |
| 23724109 | DNS resolution failed for one or more specified servers. |

| Error Code | Description |
| --- | --- |
| 23724112 | DNS resolution failed due to an internal error. Contact technical support if this issue persists. |
| 23724132 | DNS resolution failed for all the specified servers. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces |

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

cluster_nis_service

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create the NIS configuration for the cluster**

POST `/security/authentication/cluster/nis`

**Introduced In:** 9.6

The cluster can have one NIS server configuration. Specify the NIS domain and NIS servers as input. The servers field cannot be empty.
Both FQDNs and IP addresses are supported for the `server` property. IPv6 must be enabled if IPv6 family addresses are specified in the `server` property. A maximum of ten NIS servers are supported.

**Required properties**

- `domain` - NIS domain to which this configuration belongs.
- `servers` - List of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

**Request Body**

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |

| Name | Type | Description |
| --- | --- | --- |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used<br>by the NIS domain configuration. |

**Example request**

```
{
  "binding_details": [
    {
      "server": "string",
      "status": {
        "code": "string",
        "message": "string"
      }
    }
  ],
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of NIS domain records. |
| records | array[cluster_nis_service] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "binding_details": [
        {
          "server": "string",
          "status": {
            "code": "string",
            "message": "string"
          }
        }
      ],
      "bound_servers": [
        "string"
      ],
      "domain": "string",
      "servers": [
        "string"
      ]
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1966253 | IPv6 is not enabled in the cluster. |

| Error Code | Description |
| --- | --- |
| 3276964 | The NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters. |
| 3276933 | A maximum of 10 NIS servers can be configured per SVM. |
| 13434916 | The SVM is in the process of being created. Wait a few minutes, and then try the command again. |
| 23724109 | DNS resolution failed for one or more specified servers. |
| 23724112 | DNS resolution failed due to an internal error. Contact technical support if this issue persists. |
| 23724132 | DNS resolution failed for all the specified servers. |
| 23724130 | Cannot use an IPv6 name server address because there are no IPv6 interfaces. |
| 23724111 | Invalid value specified for nis-servers. The value must be a valid hostname or IP address. |
| 23724087 | The specified IPv6 address is not supported because it is one of the following: ::, link-local, multicast, v4-compatible, v4-mapped, loopback. |
| 23724086 | The specified IPv4 address is not supported because it is one of the following: multicast, loopback, 0.0.0.0 or broadcast. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

binding_status

| Name | Type | Description |
|---|---|---|
| code | string | Code corresponding to the server's binding status. |
| message | string | Detailed description of the server's binding status. |

binding_details

| Name | Type | Description |
|---|---|---|
| server | string | Hostname/IP address of the NIS server in the domain. |
| status | binding_status | |

cluster_nis_service

| Name | Type | Description |
|---|---|---|
| binding_details | array[binding_details] | An array of objects where each object represents the NIS server and it's status for a given NIS domain. It is an advanced field. |
| bound_servers | array[string] | |
| domain | string | The NIS domain to which this configuration belongs. |
| servers | array[string] | A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and update OAuth 2.0 status

### View and update OAuth 2.0 status

#### Overview

This API is used to enable and disable OAuth 2.0 in the cluster. The PATCH request enables and disables OAuth 2.0 in the cluster.
+

---

#### Examples

### Retrieving the OAuth 2.0 status in the cluster

The following output shows the OAuth 2.0 status of the cluster.
+

---

```
# The API:
/api/security/authentication/cluster/oauth2

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/oauth2"
-H "accept: application/hal+json"

# The response:
{
"enabled": false
}
```

---

#### Enabling the OAuth 2.0 in the cluster

The following output shows how to enable the OAuth 2.0 in the cluster.
+

---

```
# The API:
/api/security/authentication/cluster/oauth2

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/authentication/cluster/oauth2" -H "accept:
application/hal+json" -d '{ "enabled": true }'
```

---

### Retrieve the OAuth 2.0 status

GET `/security/authentication/cluster/oauth2`

**Introduced In:** 9.14

Retrieves the OAuth 2.0 status.

#### Related ONTAP commands

- `security oauth2 show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| enabled | boolean | query | False | Filter by enabled |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether OAuth 2.0 is enabled or disabled globally. Default value is false. |

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the OAuth 2.0 status

PATCH `/security/authentication/cluster/oauth2`

**Introduced In:** 9.14

Updates the OAuth 2.0 status.

#### Related ONTAP commands

- `security oauth2 modify`

#### Request Body

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether OAuth 2.0 is enabled or disabled globally. Default value is false. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default,
```

#### Definitions

**See Definitions**

security_oauth2_global

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Indicates whether OAuth 2.0 is enabled or disabled globally. Default value is false. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and create OAuth 2.0 configurations

### View and create OAuth 2.0 configurations

#### Overview

This API is used to retrieve and configure relevant information pertaining to the OAuth 2.0 configuration in the cluster. The POST request creates the OAuth 2.0 configuration if there is none present. Various responses are shown in the examples below.
+

**Examples**

### Retrieving the OAuth 2.0 configuration in the cluster

The following output shows the OAuth 2.0 configuration in the cluster.
+

```
# The API:
/api/security/authentication/cluster/oauth2/clients

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/authentication/cluster/oauth2/clients" -H "accept:
application/hal+json"

# The response:
{
 "records": [
{
  "name": "auth0"
}
 ],
 "num_records": 1
}
```

### Creating the OAuth 2.0 configuration

The following output shows how to create the OAuth 2.0 configuration in the cluster.
+

```
= The API:

/api/security/authentication/cluster/oauth2/clients

= The call:

curl -X POST "https://+++<mgmt-
ip>+++/api/security/authentication/cluster/oauth2/clients?return_records=t
rue" -H "accept: application/hal+json" -d '{ "name": "name",
"application": "http", "issuer": "https://examplelab.customer.com",
"audience": "aud", "client_id": "client_id", "client_secret":
```

```
"client_secret", "introspection": {"endpoint_uri":
"https://examplelab.customer.com/server/endpoint", "interval": "PT1H" },
"remote_user_claim": "user_claim", "outgoing_proxy":
"https://johndoe:somesecret@proxy.example.com:8080",
"use_local_roles_if_present": false, "use_mutual_tls": "required"
}'+++</mgmt-ip>+++

= The response:

{
"job": {
  "uuid": "e45b123b-c228-11e8-aa20-0050568e36bb",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e45b123b-c228-11e8-aa20-0050568e36bb"
    }
  }
}
}

'''





[[IDcf298733148ca2b653d4f6d54fe7d505]]
= Retrieve all OAuth 2.0 configurations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/cluster/oauth2/clients`#

*Introduced In:* 9.14

Retrieves all OAuth 2.0 configurations.

== Related ONTAP commands

* `security oauth2 client show`


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|issuer
|string
|query
|False
a|Filter by issuer


|remote_user_claim
|string
|query
|False
a|Filter by remote_user_claim


|use_local_roles_if_present
|boolean
|query
|False
a|Filter by use_local_roles_if_present


|introspection.endpoint_uri
|string
|query
|False
a|Filter by introspection.endpoint_uri


|introspection.interval
|string
|query
|False
a|Filter by introspection.interval


|client_id
|string
|query
|False
a|Filter by client_id
```

```
|use_mutual_tls
|string
|query
|False
a|Filter by use_mutual_tls


|provider
|string
|query
|False
a|Filter by provider

* Introduced in: 9.16


|name
|string
|query
|False
a|Filter by name


|application
|string
|query
|False
a|Filter by application


|audience
|string
|query
|False
a|Filter by audience


|jwks.provider_uri
|string
|query
|False
a|Filter by jwks.provider_uri


|jwks.refresh_interval
|string
|query
```

```
|False
a|Filter by jwks.refresh_interval


|hashed_client_secret
|string
|query
|False
a|Filter by hashed_client_secret


|outgoing_proxy
|string
|query
|False
a|Filter by outgoing_proxy


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_oauth2[security_oauth2]]
a|

|===


.Example response
[%collapsible%closed]
```

```
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": "string",
      "audience": "string",
      "client_id": "string",
      "hashed_client_secret": "string",
      "introspection": {
        "endpoint_uri":
"https://examplelab.customer.com/token/introspect",
        "interval": "PT1H"
      },
      "issuer": "https://examplelab.customer.com",
      "jwks": {
        "provider_uri": "https://examplelab.customer.com/pf/JWKS",
        "refresh_interval": "PT2H"
      },
      "name": "auth0",
      "outgoing_proxy":
"https://johndoe:secretpass@proxy.example.com:8080",
      "provider": "entra",
      "remote_user_claim": "string",
      "skip_uri_validation": true,
      "use_mutual_tls": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#introspection]
[.api-collapsible-fifth-title]
introspection

[cols=3*,options=header]
|===
|Name
|Type
|Description

|endpoint_uri
|string
a|The token introspection endpoint URI.


|interval
|string
a|The refresh interval for caching tokens, in ISO-8601 format. This can be
set to the value "disabled" to disable caching of tokens. When set to 0,
tokens are cached according to the expiry period in them. Otherwise, it
can be set to a value from 1 second to 2147483647 seconds.


|===


[#jwks]
[.api-collapsible-fifth-title]
jwks

[cols=3*,options=header]
|===
|Name
|Type
|Description

|provider_uri
|string
a|The URI on which the JSON Web Key Set (JWKS) are hosted.


|refresh_interval
|string
a|The refresh interval for the JSON Web Key Set (JWKS), in ISO-8601
```

```
format. This can be set to a value from 300 seconds to 2147483647 seconds.


|===


[#security_oauth2]
[.api-collapsible-fifth-title]
security_oauth2

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application
|string
a|The name of the application using OAuth 2.0. Required for POST
operations.


|audience
|string
a|The OAuth 2.0 Audience.


|client_id
|string
a|The OAuth 2.0 client ID. Required in POST operations for remote
introspection.


|hashed_client_secret
|string
a|The OAuth 2.0 client secret as a SHA256 HMAC hashed value created with
the cluster UUID as its HMAC secret key.


|introspection
|link:#introspection[introspection]
a|
```

```
|issuer
|string
a|The OAuth 2.0 Issuer.


|jwks
|link:#jwks[jwks]
a|

|name
|string
a|The configuration name. Required for POST operations.


|outgoing_proxy
|string
a|Outgoing proxy to access external identity providers (IdPs). If not
specified, no proxy is configured.


|provider
|string
a|The Identity Provider type.


|remote_user_claim
|string
a|The remote user claim.


|use_local_roles_if_present
|boolean
a|Indicates whether or not to use locally configured roles, if present.
Default value is false.


|use_mutual_tls
|string
a|OAuth 2.0 mutual TLS authentication setting. Set this value to "none" to
disable mutual TLS authentication. Set this value to "required" to enforce
mutual TLS authentication for all access tokens and reject any token that
does not have x5t#S256 property in the cnf section. The default value is
"request" which means mutual TLS authentication is enforced only if the
x5t#S256 property is present in the cnf section of the access token.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID23b7dfc2a7af337bf407df58ec2d7a78]]
= Create the OAuth 2.0 configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/cluster/oauth2/clients`#

*Introduced In:* 9.14

Creates the OAuth 2.0 configuration.

== Required properties

* `name`
* `application`
* `issuer`

== Optional properties

* `audience`
* `client_id`
* `client_secret`
* `introspection.endpoint_uri`
* `introspection.interval`
* `remote_user_claim`
* `jwks.provider_uri`
* `jwks.refresh_interval`
* `outgoing_proxy`
* `use_local_roles_if_present`
* `skip_uri_validation`
* `use_mutual_tls`
* `provider`
```

```
== Related ONTAP commands

* `security oauth2 client create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|application
|string
a|The name of the application using OAuth 2.0. Required for POST
operations.


|audience
|string
a|The OAuth 2.0 Audience.


|client_id
|string
a|The OAuth 2.0 client ID. Required in POST operations for remote
introspection.


|client_secret
|string
a|The OAuth 2.0 client secret. Required in POST operations for remote
introspection.


|hashed_client_secret
|string
a|The OAuth 2.0 client secret as a SHA256 HMAC hashed value created with
the cluster UUID as its HMAC secret key.


|introspection
|link:#introspection[introspection]
a|

|issuer
|string
a|The OAuth 2.0 Issuer.


|jwks
|link:#jwks[jwks]
```

```
a|

|name
|string
a|The configuration name. Required for POST operations.


|outgoing_proxy
|string
a|Outgoing proxy to access external identity providers (IdPs). If not
specified, no proxy is configured.


|provider
|string
a|The Identity Provider type.


|remote_user_claim
|string
a|The remote user claim.


|skip_uri_validation
|boolean
a|Indicates whether or not to validate the input URIs. Default value is
false.


|use_local_roles_if_present
|boolean
a|Indicates whether or not to use locally configured roles, if present.
Default value is false.


|use_mutual_tls
|string
a|OAuth 2.0 mutual TLS authentication setting. Set this value to "none" to
disable mutual TLS authentication. Set this value to "required" to enforce
mutual TLS authentication for all access tokens and reject any token that
does not have x5t#S256 property in the cnf section. The default value is
"request" which means mutual TLS authentication is enforced only if the
x5t#S256 property is present in the cnf section of the access token.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "application": "string",
  "audience": "string",
  "client_id": "string",
  "client_secret": "string",
  "hashed_client_secret": "string",
  "introspection": {
    "endpoint_uri": "https://examplelab.customer.com/token/introspect",
    "interval": "PT1H"
  },
  "issuer": "https://examplelab.customer.com",
  "jwks": {
    "provider_uri": "https://examplelab.customer.com/pf/JWKS",
    "refresh_interval": "PT2H"
  },
  "name": "auth0",
  "outgoing_proxy": "https://johndoe:secretpass@proxy.example.com:8080",
  "provider": "entra",
  "remote_user_claim": "string",
  "use_mutual_tls": "string"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 203817010
| Client ID is required for remote introspection.

| 203817011
| Client secret is required for remote introspection.

| 203817012
| Client ID and client secret required for remote introspection.

| 203817013
| JWKS URI should not be configured for remote introspection.

| 203817014
| JWKS refresh interval should not be specified for remote introspection.

| 203817015
| The token introspection endpoint is required for remote introspection.

| 203817016
| JWKS refresh interval provided without providing JWKS URI.

| 203817017
| Minimum supported value of JWKS refresh interval is 300 seconds.

| 203817018
| Required parameters for either local validation or remote introspection
are missing. Provide either the JWKS URI for local validation, or metadata
configuration URI or token introspection endpoint with client ID and
secret for remote introspection.

| 203817019
| Failed to add new IDP client because number of maximum supported IDP
clients is already reached.

| 203817020
| Internal error. Failed to validate provider URI.
```

```
| 203817021
| OAuth 2.0 Provider URI validation failed with error.

| 203817022
| OAuth 2.0 Provider JWKS URI validation failed. Received empty response
message from the JWKS URI.

| 203817023
| OAuth 2.0 Provider JWKS URI validation failed. No keys were found in
response message received from the JWKS URI.

| 203817024
| Internal error. Unable to allocate memory for CURL response.

| 203817025
| Maximum value of JWKS refresh interval is 2147483647 seconds.

| 203817033
| OAuth 2.0 Provider Introspection endpoint validation failed. Received
empty response message from the Introspection endpoint.

| 203817034
| OAuth 2.0 Provider Introspection endpoint validation failed. Received
invalid response message for Introspection request.

| 203817037
| An entry cannot be created as another entry with the same configuration
name has the same issuer and audience.

| 203817042
| Maximum value of introspection interval is 2147483647 seconds.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
```

```
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#introspection]
[.api-collapsible-fifth-title]
introspection

[cols=3*,options=header]
|===
|Name
|Type
|Description

|endpoint_uri
|string
a|The token introspection endpoint URI.


|interval
|string
a|The refresh interval for caching tokens, in ISO-8601 format. This can be
set to the value "disabled" to disable caching of tokens. When set to 0,
tokens are cached according to the expiry period in them. Otherwise, it
can be set to a value from 1 second to 2147483647 seconds.


|===


[#jwks]
[.api-collapsible-fifth-title]
```

```
jwks

[cols=3*,options=header]
|===
|Name
|Type
|Description

|provider_uri
|string
a|The URI on which the JSON Web Key Set (JWKS) are hosted.


|refresh_interval
|string
a|The refresh interval for the JSON Web Key Set (JWKS), in ISO-8601
format. This can be set to a value from 300 seconds to 2147483647 seconds.


|===


[#security_oauth2]
[.api-collapsible-fifth-title]
security_oauth2

[cols=3*,options=header]
|===
|Name
|Type
|Description

|application
|string
a|The name of the application using OAuth 2.0. Required for POST
operations.


|audience
|string
a|The OAuth 2.0 Audience.


|client_id
|string
a|The OAuth 2.0 client ID. Required in POST operations for remote
introspection.
```

```
|client_secret
|string
a|The OAuth 2.0 client secret. Required in POST operations for remote
introspection.


|hashed_client_secret
|string
a|The OAuth 2.0 client secret as a SHA256 HMAC hashed value created with
the cluster UUID as its HMAC secret key.


|introspection
|link:#introspection[introspection]
a|

|issuer
|string
a|The OAuth 2.0 Issuer.


|jwks
|link:#jwks[jwks]
a|

|name
|string
a|The configuration name. Required for POST operations.


|outgoing_proxy
|string
a|Outgoing proxy to access external identity providers (IdPs). If not
specified, no proxy is configured.


|provider
|string
a|The Identity Provider type.


|remote_user_claim
|string
a|The remote user claim.
```

```
|skip_uri_validation
|boolean
a|Indicates whether or not to validate the input URIs. Default value is
false.


|use_local_roles_if_present
|boolean
a|Indicates whether or not to use locally configured roles, if present.
Default value is false.


|use_mutual_tls
|string
a|OAuth 2.0 mutual TLS authentication setting. Set this value to "none" to
disable mutual TLS authentication. Set this value to "required" to enforce
mutual TLS authentication for all access tokens and reject any token that
does not have x5t#S256 property in the cnf section. The default value is
"request" which means mutual TLS authentication is enforced only if the
x5t#S256 property is present in the cnf section of the access token.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and delete an OAuth 2.0 configuration

:leveloffset: +1


[[ID66bd187c30364ece541fccbc867dc0c9]]
= View and delete an OAuth 2.0 configuration



== Overview

This API is used to retrieve and delete the OAuth 2.0 configuration in the
cluster. The GET request retrieves the OAuth 2.0 configuration. The DELETE
request removes the OAuth 2.0 configuration.  Various responses are shown
in the examples below.
 +

'''

== Examples

=== Retrieving the OAuth 2.0 configuration in the cluster

The following output shows the OAuth 2.0 configuration in the cluster.
 +

'''

----

# The API:
```

```
/api/security/authentication/cluster/oauth2/clients/{name}

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/authentication/cluster/oauth2/clients/auth0" -H "accept:
application/hal+json"

# The response:
{
"name": "auth0",
"application": "http",
"issuer": "https://examplelab.customer.com",
"audience": "aud",
"client_id": "client_id",
"hashed_client_secret": "<HASHED-CLIENT-SECRET>",
"introspection": {
  "endpoint_uri": "https://examplelab.customer.com/server/endpoint",
  "interval": "PT1H"
},
"remote_user_claim": "user_claim",
"jwks": {
  "provider_uri": "https://examplelab.customer.com/pf/JWKS",
  "refresh_interval": "PT1H"
},
"use_local_roles_if_present": false,
"outgoing_proxy": "https://johndoe:secretpass@proxy.example.com:8080",
"_links": {
  "self": {
    "href": "/api/security/authentication/cluster/oauth2/clients"
  }
},
"use_mutual_tls": "required"
}
----

'''

=== Deleting the OAuth 2.0 configuration

'''

----

# The API:
/api/security/authentication/cluster/oauth2/clients/{name}
```

```
# The call:
curl -X DELETE "https://<mgmt-
ip>/api/security/authentication/cluster/oauth2/clients/auth0"
----

'''




[[ID0cdde8ff2e54a8b9c29b333e3dd48b8e]]
= Delete an OAuth 2.0 configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/cluster/oauth2/clients/{name}`#

*Introduced In:* 9.14

Deletes the OAuth 2.0 configuration with the specified name.

== Required properties

***** `config_name`

== Related ONTAP commands

* `security oauth2 client delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|OAuth 2.0 configuration name.
```

```
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 203816995
| OAuth 2.0 must be disabled before the configuration can be removed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6baef89d7e26dfb77463723860dd6596]]
= Retrieve an OAuth 2.0 configuration with the specified name

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/cluster/oauth2/clients/{name}`#

*Introduced In:* 9.14

Retrieves the OAuth 2.0 configuration with the specified name.

== Related ONTAP commands

* `security oauth2 client show`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|OAuth 2.0 configuration name.


|issuer
|string
|query
|False
a|Filter by issuer


|remote_user_claim
|string
|query
|False
a|Filter by remote_user_claim


|use_local_roles_if_present
|boolean
|query
|False
a|Filter by use_local_roles_if_present


|introspection.endpoint_uri
|string
|query
|False
a|Filter by introspection.endpoint_uri
```

```
|introspection.interval
|string
|query
|False
a|Filter by introspection.interval


|client_id
|string
|query
|False
a|Filter by client_id


|use_mutual_tls
|string
|query
|False
a|Filter by use_mutual_tls


|provider
|string
|query
|False
a|Filter by provider

* Introduced in: 9.16


|application
|string
|query
|False
a|Filter by application


|audience
|string
|query
|False
a|Filter by audience


|jwks.provider_uri
|string
|query
```

```
|False
a|Filter by jwks.provider_uri


|jwks.refresh_interval
|string
|query
|False
a|Filter by jwks.refresh_interval


|hashed_client_secret
|string
|query
|False
a|Filter by hashed_client_secret


|outgoing_proxy
|string
|query
|False
a|Filter by outgoing_proxy


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|application
|string
a|The name of the application using OAuth 2.0. Required for POST
operations.


|audience
|string
a|The OAuth 2.0 Audience.


|client_id
|string
a|The OAuth 2.0 client ID. Required in POST operations for remote
introspection.


|hashed_client_secret
|string
a|The OAuth 2.0 client secret as a SHA256 HMAC hashed value created with
the cluster UUID as its HMAC secret key.


|introspection
|link:#introspection[introspection]
a|

|issuer
|string
a|The OAuth 2.0 Issuer.


|jwks
|link:#jwks[jwks]
a|

|name
|string
a|The configuration name. Required for POST operations.


|outgoing_proxy
|string
a|Outgoing proxy to access external identity providers (IdPs). If not
specified, no proxy is configured.
```

```
|provider
|string
a|The Identity Provider type.


|remote_user_claim
|string
a|The remote user claim.


|use_local_roles_if_present
|boolean
a|Indicates whether or not to use locally configured roles, if present.
Default value is false.


|use_mutual_tls
|string
a|OAuth 2.0 mutual TLS authentication setting. Set this value to "none" to
disable mutual TLS authentication. Set this value to "required" to enforce
mutual TLS authentication for all access tokens and reject any token that
does not have x5t#S256 property in the cnf section. The default value is
"request" which means mutual TLS authentication is enforced only if the
x5t#S256 property is present in the cnf section of the access token.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application": "string",
  "audience": "string",
  "client_id": "string",
  "hashed_client_secret": "string",
  "introspection": {
    "endpoint_uri": "https://examplelab.customer.com/token/introspect",
    "interval": "PT1H"
```

```
  },
  "issuer": "https://examplelab.customer.com",
  "jwks": {
    "provider_uri": "https://examplelab.customer.com/pf/JWKS",
    "refresh_interval": "PT2H"
  },
  "name": "auth0",
  "outgoing_proxy": "https://johndoe:secretpass@proxy.example.com:8080",
  "provider": "entra",
  "remote_user_claim": "string",
  "skip_uri_validation": true,
  "use_mutual_tls": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#introspection]
[.api-collapsible-fifth-title]
introspection

[cols=3*,options=header]
|===
|Name
|Type
|Description

|endpoint_uri
|string
a|The token introspection endpoint URI.


|interval
|string
a|The refresh interval for caching tokens, in ISO-8601 format. This can be
set to the value "disabled" to disable caching of tokens. When set to 0,
tokens are cached according to the expiry period in them. Otherwise, it
can be set to a value from 1 second to 2147483647 seconds.


|===


[#jwks]
[.api-collapsible-fifth-title]
jwks

[cols=3*,options=header]
|===
|Name
|Type
|Description

|provider_uri
|string
a|The URI on which the JSON Web Key Set (JWKS) are hosted.


|refresh_interval
|string
a|The refresh interval for the JSON Web Key Set (JWKS), in ISO-8601
format. This can be set to a value from 300 seconds to 2147483647 seconds.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SAML service

:leveloffset: +1


[[IDb7b9e547d7200fa113009944a544d1f6]]
= Manage SAML service



== Overview

This API is used to retrieve and display relevant information pertaining
to the SAML service provider configuration in the cluster. The POST
request creates a SAML service provider configuration if there is none
present.  The DELETE request removes the SAML service provider
configuration.  The PATCH request enables and disables SAML in the
cluster.  Various responses are shown in the examples below.
 +

'''

== Examples

=== Retrieving the SAML service provider configuration in the cluster

The following output shows the SAML service provider configuration in the
```

```
cluster.
 +

'''

----

# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/saml-
sp" -H "accept: application/hal+json"

# The response:
{
"idp_uri": "https://examplelab.customer.com/idp/Metadata",
"enabled": true,
"host": "172.21.74.181",
"certificate": {
  "ca": "cluster1",
  "serial_number": "156F10C3EB4C51C1"
},
"_links": {
  "self": {
    "href": "/api/security/authentication/cluster/saml-sp"
  }
}
}
----

'''

=== Creating the SAML service provider configuration

The following output shows how to create a SAML service provider
configuration in the cluster.
 +

'''

----

# The API:
/api/security/authentication/cluster/saml-sp

# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/saml-
sp?return_records=true" -H "accept: application/hal+json" -d '{ "idp_uri":
"https://examplelab.customer.com/idp/Metadata", "host": "172.21.74.181",
"certificate": { "ca": "cluster1", "serial_number": "156F10C3EB4C51C1" }}'
----

'''

=== Updating the SAML service provider configuration

The following output shows how to enable a SAML service provider
configuration in the cluster.

Disabling the configuration requires the client to be authenticated
through SAML prior to performing the operation.
 +

'''

----

# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/saml-
sp/" -d '{ "enabled": true }'
----

'''

=== Deleting the SAML service provider configuration

'''

----

# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/security/authentication/cluster/saml-sp/"
----

'''
```

```
[[ID1b4294c6cd4416ff0bc13a58bdcd9d3c]]
= Delete an SAML service provider configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp`#

*Introduced In:* 9.6

Deletes a SAML service provider configuration.


== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 12320803
| SAML must be disabled before the configuration can be removed.

| 12320832
| Internal Error. Failed to execute the command, web server setup is in
progress. Retry after some time.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[ID57bce561cd311f63e44c0d995c189f29]]
= Retrieve an SAML service provider configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp`#

*Introduced In:* 9.6

Retrieves a SAML service provider configuration.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|host
|string
|query
|False
a|Filter by host

* Introduced in: 9.7


|certificate.ca
|string
|query
|False
a|Filter by certificate.ca

* Introduced in: 9.7
* maxLength: 256
* minLength: 1


|certificate.serial_number
```

```
|string
|query
|False
a|Filter by certificate.serial_number

* Introduced in: 9.7
* maxLength: 40
* minLength: 1


|certificate.common_name
|string
|query
|False
a|Filter by certificate.common_name

* Introduced in: 9.7


|enabled
|boolean
|query
|False
a|Filter by enabled

* Introduced in: 9.7


|idp_uri
|string
|query
|False
a|Filter by idp_uri

* Introduced in: 9.7


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|certificate
|link:#certificate[certificate]
a|

|enabled
|boolean
a|The SAML service provider is enabled.  Valid for PATCH and GET
operations only.


|host
|string
a|The SAML service provider host.


|idp_uri
|string
a|The identity provider (IdP) metadata location. Required for POST
operations.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "host": "string",
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-
06/FederationMetadata.xml"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#certificate]
[.api-collapsible-fifth-title]
certificate
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|ca
|string
a|Server certificate issuing certificate authority (CA).  This cannot be
used with the server certificate common name.


|common_name
|string
a|Server certificate common name.  This cannot be used with the
certificate authority (CA) or serial_number.


|serial_number
|string
a|Server certificate serial number.  This cannot be used with the server
certificate common name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID2abdacc6ba8b2db4a902e94b9fef8255]]
= Update an SAML service provider configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp`#
```

```
*Introduced In:* 9.6

Updates a SAML service provider configuration.


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|The SAML service provider is enabled.  Valid for PATCH and GET
operations only.


|===


== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 12320791
| SAML can only be disabled using the console or a SAML-authenticated
application.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate
[#security_saml_sp]
[.api-collapsible-fifth-title]
security_saml_sp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|The SAML service provider is enabled.  Valid for PATCH and GET
operations only.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0265e88436b155f199d0c78153720fb2]]
= Create an SAML service provider configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp`#

*Introduced In:* 9.6

Creates a SAML service provider configuration. Note that "common_name" is
mutually exclusive with "serial_number" and "ca" in POST. SAML will
initially be disabled, requiring a patch to set "enabled" to "true", so
that the user has time to complete the setup of the IdP.

== Required properties

* `idp_uri`

== Optional properties

* `certificate`
* `enabled`
* `host`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|verify_metadata_server
|boolean
|query
|False
a|Verify IdP metadata server identity.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|
```

```
|host
|string
a|The SAML service provider host.


|idp_uri
|string
a|The identity provider (IdP) metadata location. Required for POST
operations.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "enabled": true,
  "host": "string",
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-
06/FederationMetadata.xml"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 12320789
| Failed to download data file from specified URI.

| 12320794
| The host parameter provided must be the cluster management interface's
IP address.  If the cluster management interface is not available, the
node management interface's IP address must be used.

| 12320795
| A valid cluster or node management interface IP address must be
provided.

| 12320805
| The certificate information provided does not match any installed
certificates.

| 12320806
| The certificate information entered does not match any installed
certificates.

| 12320814
| An invalid IDP URI has been entered.

| 12320815
| An IDP URI must be an HTTPS or FTPS URI.

| 12320819
| Use the HTTPS scheme for the `idp_uri` or set `verify_metadata_server`
to false.

| 12320820
| No certificate is installed with the specified `certificate.ca` and
`certificate.serial`.

| 12320821
| No certificate is installed with the specified
```

```
`certificate.common_name`.

| 12320823
| The host parameter provided must be the cluster management interface's
IP address. If the cluster management interface is not available, the node
management interface's IP address must be used.

| 12320827
| The certificate authority specified in the default metadata
configuration does not match the provided certificate authority.

| 12320828
| The SAML service provider host specified in the default metadata
configuration does not match the provided SAML service provider host.

| 12320829
| The certificate serial number specified in the default metadata
configuration does not match the provided certificate serial number.

| 12320830
| The certificate common name specified in the default metadata
configuration does not match the provided certificate common name.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ca
|string
a|Server certificate issuing certificate authority (CA).  This cannot be
used with the server certificate common name.


|common_name
|string
a|Server certificate common name.  This cannot be used with the
certificate authority (CA) or serial_number.


|serial_number
|string
a|Server certificate serial number.  This cannot be used with the server
certificate common name.


|===


[#security_saml_sp]
[.api-collapsible-fifth-title]
security_saml_sp

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|certificate
|link:#certificate[certificate]
a|

|host
|string
a|The SAML service provider host.


|idp_uri
|string
a|The identity provider (IdP) metadata location. Required for POST
operations.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SAML default metadata configuration

:leveloffset: +1


[[IDfab22cb08c419d8db2531187c7c76cf9]]
= Manage SAML default metadata configuration



== Overview

This API is used to manage relevant information about the SAML default
metadata configuration in the cluster. The POST request creates a SAML
default metadata configuration if there is none present.  The DELETE
request removes the SAML default metadata configuration.  Various
responses are shown in the examples below.
 +

'''

== Examples

=== Retrieving the SAML default metadata configuration in the cluster

The following output shows the SAML default metadata configuration in the
cluster.
 +

'''

----

# The API:
/api/security/authentication/cluster/saml-sp/default-metadata

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/saml-
sp/default-metadata" -H "accept: application/hal+json"

# The response:
{
"host": "172.21.74.181",
"certificate": {
  "ca": "cluster-1",
  "serial_number": "180E3331A0DC5A19",
  "common_name": "cluster-1"
},
"scope": "cluster",
"_links": {
  "self": {
    "href": "/api/security/authentication/cluster/saml-sp/default-
metadata"
  }
}
}
----

'''

=== Creating the SAML default metadata configuration in the cluster

The following output shows how to create the SAML default metadata in the
cluster.
 +

'''

----

# The API:
/api/security/authentication/cluster/saml-sp/default-metadata

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/saml-
sp/default-metadata?return_records=true" -H "accept: application/hal+json"
-d '{ "host": "172.21.74.181", "certificate": { "ca": "cluster1",
"serial_number": "156F10C3EB4C51C1" }}'
----

'''

=== Deleting the SAML default metadata configuration in the cluster
```

```
'''

----

# The API:
/api/security/authentication/cluster/saml-sp/default-metadata

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/security/authentication/cluster/saml-sp/default-metadata/"
----

'''





[[ID15a9eddb6687f59d93777e277d683e33]]
= Delete the SAML default metadata configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp/default-metadata`#

*Introduced In:* 9.17

Deletes the SAML default metadata configuration.


== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 12320817
| Internal Error. The SAML service provider metadata configuration is
being created, but is taking longer than expected to complete.
```

```
| 12320826
| SAML configuration already exists. Cannot delete the metadata when SAML
configuration is already present.

| 12320832
| Internal Error. Failed to execute the command, web server setup is in
progress. Retry after some time.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf6c60b0ec9e34484fa185ccbcc4cc6e8]]
= Retrieve the SAML default metadata configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp/default-metadata`#

*Introduced In:* 9.17

Retrieves the SAML default metadata configuration.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|host
|string
```

```
|query
|False
a|Filter by host


|certificate.common_name
|string
|query
|False
a|Filter by certificate.common_name


|certificate.serial_number
|string
|query
|False
a|Filter by certificate.serial_number

* maxLength: 40
* minLength: 1


|certificate.ca
|string
|query
|False
a|Filter by certificate.ca

* maxLength: 256
* minLength: 1


|scope
|string
|query
|False
a|Filter by scope


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|certificate
|link:#certificate[certificate]
a|

|host
|string
a|The SAML service provider host.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
```

```
  },
  "host": "string",
  "scope": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#certificate]
[.api-collapsible-fifth-title]
certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ca
|string
a|Server certificate issuing certificate authority (CA).  This cannot be
used with the server certificate common name.


|common_name
|string
a|Server certificate common name.  This cannot be used with the
certificate authority (CA) or serial_number.


|serial_number
|string
a|Server certificate serial number.  This cannot be used with the server
certificate common name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID485be2866f13ec89984085534fc90db3]]
= Create the SAML default metadata configuration


[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/cluster/saml-sp/default-metadata`#

*Introduced In:* 9.17

Creates the SAML default metadata configuration. Note that `common_name`
is mutually exclusive with `serial_number` and `ca` in POST requests.
```

```
== Optional properties

* `certificate`
* `host`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
```

```
a|

|host
|string
a|The SAML service provider host.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "host": "string",
  "scope": "string"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 12320794
| Invalid host parameter for default metadata generation. The host
parameter must be the cluster management interface's IP address. If the
cluster management interface is not available, the node management
interface's IP address must be used.

| 12320805
| The certificate information provided does not match any installed
certificates.

| 12320806
| The certificate information entered does not match any installed
certificates.

| 12320820
| No certificate is installed with the specified `certificate.ca` and
`certificate.serial`.

| 12320821
| No certificate is installed with the specified
`certificate.common_name`.

| 12320825
| SAML configuration already exists. Cannot create the metadata when SAML
configuration is already present.

| 12320831
| Default metadata configuration already exists.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ca
|string
a|Server certificate issuing certificate authority (CA).  This cannot be
used with the server certificate common name.


|common_name
|string
a|Server certificate common name.  This cannot be used with the
certificate authority (CA) or serial_number.
```

```
|serial_number
|string
a|Server certificate serial number.  This cannot be used with the server
certificate common name.


|===


[#security_saml_def_metadata]
[.api-collapsible-fifth-title]
security_saml_def_metadata

[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|

|host
|string
a|The SAML service provider host.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and create Duo groups

:leveloffset: +1


[[ID8cfdc5adc4352b5de13662570422cb6d]]
= View and create Duo groups



== Overview

This API configures the group to be included in Duo authentication.

== Examples

=== Creating a Duo group

Specify the group name and other necessary information in the body of the
POST request. The owner.uuid or owner.name are not required for a cluster-
```

```
scoped account.

----

# The API:
POST "/api/security/authentication/duo/groups"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/duo/groups --request
POST --data '{"name": "test",  "comment": "test post"}'

# The response:
{
}
----

=== Creating a Duo group for SVM-scoped account

For an SVM-scoped account, specify either the SVM name as the owner.name
or the SVM UUID as the owner.uuid along with other parameters for the
group. These parameters indicate the SVM that contains the user group and
can be obtained from the response body of the GET request performed on the
API "/api/svm/svms".

----

# The API:
POST "/api/security/authentication/duo/groups"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/duo/groups --request
POST --data '{"name": "test",  "comment": "test post", "owner.name":"vs0"
}'

# The response:
{
}
----

=== Retrieving the configured Duo groups

Retrieves the Duo groups associated with the owners or filtered list (for
a specific SVM and so on).

----

# The API:
```

```
GET "/api/security/authentication/duo/groups"

# The call to retrieve Duo groups with the cluster or SVM:
curl -k https://<mgmt-ip>/api/security/authentication/duo/groups

# the response:
{
"records": [
{
  "owner": {
  "uuid": "5f618c80-d90a-11ed-a6e6-0050568e8ef2",
  "name": "vs0"
}
},
{
"owner": {
"uuid": "f810005a-d908-11ed-a6e6-0050568e8ef2",
"name": "cluster-1"
}
}
],
"num_records": 2
}
----




[[ID32e84c74b6a057287af48a7c807dff35]]
= Retrieve the configured groups

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/duo/groups`#

*Introduced In:* 9.14

Retrieves the configured groups.

== Related ONTAP commands

* `security login duo group show`

== Learn more

*
xref:{relative_path}security_authentication_duo_groups_endpoint_overview.h
```

```
tml[DOC /security/authentication/duo/groups]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|excluded_users
|string
|query
|False
a|Filter by excluded_users


|name
|string
|query
|False
a|Filter by name


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|comment
|string
|query
```

```
|False
a|Filter by comment


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
```

```
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#duogroup[duogroup]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "excluded_users": [
        "user1",
        "user2"
      ],
      "name": "AD_Group",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo group profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duogroup]
[.api-collapsible-fifth-title]
duogroup

Group profile to include in Duo authentication.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|name
|string
a|Name of the group to be included in Duo authentication.
```

```
|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo group profile is configured.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0ccff838a85d7d3ca10b833a3523c626]]
= Create a Duo group

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/duo/groups`#

*Introduced In:* 9.14

Creates a Duo group.

== Required properties

* `name` - Group name

== Related ONTAP commands

* `security login duo group create`

== Learn more

*
xref:{relative_path}security_authentication_duo_groups_endpoint_overview.h
tml[DOC /security/authentication/duo/groups]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
```

```
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|name
|string
a|Name of the group to be included in Duo authentication.
```

```
|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo group profile is configured.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "excluded_users": [
    "user1",
    "user2"
  ],
  "name": "AD_Group",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====
```

```
== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#duogroup[duogroup]]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "comment": "string",
      "excluded_users": [
        "user1",
        "user2"
      ],
      "name": "AD_Group",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo group profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duogroup]
[.api-collapsible-fifth-title]
duogroup

Group profile to include in Duo authentication.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|name
|string
a|Name of the group to be included in Duo authentication.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo group profile is configured.


|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage Duo groups configured for an SVM or cluster

:leveloffset: +1


[[IDd63b7f187861af889633854a5c7843d8]]
= Delete a Duo group configured for an SVM or cluster

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/duo/groups/{owner.uuid}/{name}`#

*Introduced In:* 9.14

Deletes a Duo group of the SVM or cluster.

== Related ONTAP commands

* `security login duo group delete`

== Learn more

*
xref:{relative_path}security_authentication_duo_groups_owner.uuid_name_end
point_overview.htmll[DOC
/security/authentication/duo/groups/{owner.uuid}/\{name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.


|name
|string
|path
|True
a|Group name.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDcc6979b77a89bbcb97c9b3017c6b0fce]]
= Retrieve the Duo group configured for an SVM or cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/duo/groups/{owner.uuid}/{name}`#
```

```
*Introduced In:* 9.14

Retrieves the Duo group configured for an SVM or cluster.

== Related ONTAP commands

* `security login duo group show`

== Learn more

*
xref:{relative_path}security_authentication_duo_groups_owner.uuid_name_end
point_overview.htmll[DOC
/security/authentication/duo/groups/{owner.uuid}/\{name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.


|name
|string
|path
|True
a|Group name.


|fields
|array[string]
|query
|False
```

```
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|name
|string
a|Name of the group to be included in Duo authentication.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo group profile is configured.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "excluded_users": [
    "user1",
    "user2"
  ],
  "name": "AD_Group",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo group profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



[[IDcc428d436c70f1ff242722f4c5903f3d]]
= Update a Duo group configured for an SVM or cluster

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/authentication/duo/groups/{owner.uuid}/{name}`#

*Introduced In:* 9.14

Updates a configured Duo group for a cluster or SVM.

== Related ONTAP commands

* `security login duo group modify`

== Learn more

*
xref:{relative_path}security_authentication_duo_groups_owner.uuid_name_end
point_overview.html[DOC
/security/authentication/duo/groups/{owner.uuid}/\{name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID
```

```
|name
|string
|path
|True
a|Group name.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "excluded_users": [
    "user1",
    "user2"
  ]
}
====
```

== Response

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo group profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duogroup]
[.api-collapsible-fifth-title]
duogroup

Group profile to include in Duo authentication.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment for the Duo group.


|excluded_users
|array[string]
a|List of excluded users.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and create a Duo profile

:leveloffset: +1


[[ID83b3d86e044f0993ffe70fab8c0c9135]]
= View and create a Duo profile



== Overview

This API configures the Duo profile for secure shell (SSH).

== Examples

=== Creating a Duo profile

Specify the integration key, secret key, API host, and other necessary
information in the body of the POST request. The owner.uuid or owner.name
are not required for a cluster-scoped account.

----

# The API:
POST "/api/security/authentication/duo/profiles"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/duo/profiles
--request POST --data '{"api_host": "api-******.duosecurity.com",
"comment": "test post",  "fail_mode": "safe",  "integration_key":
"<INTEGRATION-KEY>", "secret_key": "<SECRET-KEY>"}'

# The response:
{
}
```

```
----

=== Creating a Duo profile for an SVM-scoped account

For an SVM-scoped account, specify either the SVM name as the owner.name
or the SVM UUID as the owner.uuid along with other parameters for the Duo
profile. These parameters indicate the SVM that the Duo profile is being
created on and can be obtained from the response body of the GET request
performed on the API "/api/svm/svms".

----

# The API:
POST "/api/security/authentication/duo/profiles"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/duo/profiles
--request POST --data '{"api_host": "api-******.duosecurity.com",
"comment": "test post",  "fail_mode": "safe",  "integration_key":
"<INTEGRATION-KEY>", "secret_key": "<SECRET-KEY>", "owner.name":"vs0" }'

# The response:
{
}
----

=== Retrieving the configured Duo profile(s)

Retrieves the Duo profiles associated with the owners or filtered list
(for a specific SVM and so on).

----

# The API:
GET "/api/security/authentication/duo/profiles"

# The call to retrieve Duo profiles associated with Duo configured cluster
or SVM:
curl -k https://<mgmt-ip>/api/security/authentication/duo/profiles

# the response:
{
"records": [
{
  "owner": {
  "uuid": "5f618c80-d90a-11ed-a6e6-0050568e8ef2",
  "name": "vs0"
```

```
}
},
{
"owner": {
"uuid": "f810005a-d908-11ed-a6e6-0050568e8ef2",
"name": "tsmithcluster-1"
}
}
],
"num_records": 2
}
```


---

[[ID720c74366d48b7212cb6a0b467a2fbc9]]
= Retrieve configured Duo profiles

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/security/authentication/duo/profiles`#

*Introduced In:* 9.14

Retrieves the configured Duo profiles.

== Related ONTAP commands

* `security login duo show`

== Learn more

* xref:{relative_path}security_authentication_duo_profiles_endpoint_overview.html[DOC /security/authentication/duo/profiles]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC /security/accounts]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type

```
|In
|Required
|Description

|api_host
|string
|query
|False
a|Filter by api_host


|auto_push
|boolean
|query
|False
a|Filter by auto_push


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|integration_key
|string
|query
|False
a|Filter by integration_key


|comment
|string
|query
|False
a|Filter by comment


|fail_mode
```

```
|string
|query
|False
a|Filter by fail_mode


|fingerprint
|string
|query
|False
a|Filter by fingerprint


|max_prompts
|integer
|query
|False
a|Filter by max_prompts

* Max value: 3
* Min value: 1


|is_enabled
|boolean
|query
|False
a|Filter by is_enabled


|status
|string
|query
|False
a|Filter by status


|push_info
|boolean
|query
|False
a|Filter by push_info


|http_proxy
|string
|query
```

```
|False
a|Filter by http_proxy


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
```

```
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#duo[duo]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "api_host": "api-****.duo.com",
      "auto_push": 1,
      "comment": "string",
      "fail_mode": "safe",
      "fingerprint": "string",
      "http_proxy": "IPaddress:port",
      "integration_key": "string",
      "is_enabled": 1,
      "max_prompts": 1,
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "push_info": 1,
      "status": "OK"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duo]
[.api-collapsible-fifth-title]
duo

Duo profile for the SVM or cluster-management server (Cserver).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|api_host
|string
a|The URL at which the Duo API is hosted.


|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.
```

```
|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
|string
a|The Integration Key associated with the Duo profile.


|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo profile is configured.


|push_info
|boolean
a|Additional information sent along with the push notification for Duo
```

```
authentication.


|status
|string
a|Information on the reachability status of Duo.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDbe72a8820bc4c8f8d8d83a8748b0f09f]]
= Create a Duo profile

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/duo/profiles`#

*Introduced In:* 9.14

Creates a Duo profile.

== Required properties

* `api_host` - Duo API host
* `integration_key` - Integration key
* `secret_key` - Secret key

== Related ONTAP commands

* `security login duo create`

== Learn more
```

```
*
xref:{relative_path}security_authentication_duo_profiles_endpoint_overview
.html[DOC /security/authentication/duo/profiles]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|api_host
|string
a|The URL at which the Duo API is hosted.


|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.
```

```
|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
|string
a|The Integration Key associated with the Duo profile.


|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo profile is configured.


|push_info
|boolean
```

```
a|Additional information sent along with the push notification for Duo
authentication.


|secret_key
|string
a|The Secret Key associated with the Duo profile.


|status
|string
a|Information on the reachability status of Duo.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "api_host": "api-****.duo.com",
  "auto_push": 1,
  "comment": "string",
  "fail_mode": "safe",
  "fingerprint": "string",
  "http_proxy": "IPaddress:port",
  "integration_key": "string",
  "is_enabled": 1,
  "max_prompts": 1,
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "push_info": 1,
  "secret_key": "string",
  "status": "OK"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#duo[duo]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "api_host": "api-****.duo.com",
      "auto_push": 1,
      "comment": "string",
      "fail_mode": "safe",
      "fingerprint": "string",
      "http_proxy": "IPaddress:port",
      "integration_key": "string",
      "is_enabled": 1,
      "max_prompts": 1,
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "push_info": 1,
      "secret_key": "string",
      "status": "OK"
    }
  ]
}
====
```

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duo]
[.api-collapsible-fifth-title]
duo

Duo profile for the SVM or cluster-management server (Cserver).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|api_host
|string
a|The URL at which the Duo API is hosted.
```

```
|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.


|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
|string
a|The Integration Key associated with the Duo profile.


|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
```

```
a|SVM name and UUID for which the Duo profile is configured.


|push_info
|boolean
a|Additional information sent along with the push notification for Duo
authentication.


|secret_key
|string
a|The Secret Key associated with the Duo profile.


|status
|string
a|Information on the reachability status of Duo.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage Duo profiles for an SVM or cluster

:leveloffset: +1


[[ID15b5c1a4462af2b3e6ee78094ad3c899]]
= Manage Duo profiles for an SVM or cluster
```

```
== Overview

This API configures the Duo profile for an SVM.
Specify the owner UUID. The owner UUID corresponds to the UUID of the SVM
containing the Duo profile and can be obtained from the response body of
the GET request performed on the API "`/api/svm/svms".

== Examples

=== Retrieving the specific configured Duo profile of the cluster or SVM

----

# The API:
GET "/api/security/authentication/duo/profiles/{owner.uuid}"

# The call:
curl -k https://<mgmt-
ip>/api/security/authentication/duo/profiles/f810005a-d908-11ed-a6e6-
0050568e8ef2

# the response:
{
"owner": {
"uuid": "f810005a-d908-11ed-a6e6-0050568e8ef2",
"name": "cluster-1"
},
"integration_key": "<INTEGRATION-KEY>",
"fingerprint": "<FINGERPRINT>",
"api_host": "api-******.duosecurity.com",
"status": "Ok",
"is_enabled": true,
"auto_push": true,
"push_info": true,
"fail_mode": "safe",
"max_prompts": 1,
"comment": "Duo profile for Cserver",
}
----

=== Modifying the Duo profile

----

# The API:
```

```
PATCH "/api/security/authentication/duo/profiles/{owner.uuid}"

# The call:
curl -k "https://<mgmt-
ip>/api/security/authentication/duo/profiles/f810005a-d908-11ed-a6e6-
0050568e8ef2" --request PATCH --data "{  \"comment\": \"Testing\",
\"auto_push\": false}"

# the response:
{}
----

=== Deleting the Duo profile

----

# The API:
DELETE "/api/security/authentication/duo/profiles/{owner.uuid}"

# The call:
curl -k https://<mgmt-
ip>/api/security/authentication/duo/profiles/f810005a-d908-11ed-a6e6-
0050568e8ef2 --request DELETE

# the response:
{}
----





[[IDb13b59785988f19137092dab93fad399]]
= Delete a Duo profile for an SVM or cluster

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/duo/profiles/{owner.uuid}`#

*Introduced In:* 9.14

Deletes the Duo profile of the SVM or cluster.

== Related ONTAP commands

* `security login duo delete`

== Learn more
```

```
*
xref:{relative_path}security_authentication_duo_profiles_owner.uuid_endpoi
nt_overview.html[DOC /security/authentication/duo/profiles/{owner.uuid}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd76b3f0f5e133f0be7c5e2459c3e621f]]
= Retrieve a Duo profile configured for an SVM or cluster
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/duo/profiles/{owner.uuid}`#

*Introduced In:* 9.14

Retrieves the Duo profile configured for the cluster or an SVM.

== Related ONTAP commands

* `security login duo show`

== Learn more

*
xref:{relative_path}security_authentication_duo_profiles_owner.uuid_endpoi
nt_overview.html[DOC /security/authentication/duo/profiles/{owner.uuid}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|api_host
|string
a|The URL at which the Duo API is hosted.


|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.


|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
```

```
|string
a|The Integration Key associated with the Duo profile.


|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo profile is configured.


|push_info
|boolean
a|Additional information sent along with the push notification for Duo
authentication.


|status
|string
a|Information on the reachability status of Duo.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "api_host": "api-****.duo.com",
  "auto_push": 1,
  "comment": "string",
```

```
  "fail_mode": "safe",
  "fingerprint": "string",
  "http_proxy": "IPaddress:port",
  "integration_key": "string",
  "is_enabled": 1,
  "max_prompts": 1,
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "push_info": 1,
  "status": "OK"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
```

```
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[ID239a547802195e91043de7acd4eae7cf]]
= Update a Duo profile configured for an SVM or cluster

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/authentication/duo/profiles/{owner.uuid}`#

*Introduced In:* 9.14

Updates a configured Duo profile for a cluster or an SVM.

== Related ONTAP commands

* `security login duo modify`

== Learn more

*
xref:{relative_path}security_authentication_duo_profiles_owner.uuid_endpoi
nt_overview.html[DOC /security/authentication/duo/profiles/{owner.uuid}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|api_host
|string
a|The URL at which the Duo API is hosted.


|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.


|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
|string
a|The Integration Key associated with the Duo profile.
```

```
|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo profile is configured.


|push_info
|boolean
a|Additional information sent along with the push notification for Duo
authentication.


|secret_key
|string
a|The Secret Key associated with the Duo profile.


|status
|string
a|Information on the reachability status of Duo.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "api_host": "api-****.duo.com",
  "auto_push": 1,
  "comment": "string",
  "fail_mode": "safe",
  "fingerprint": "string",
  "http_proxy": "IPaddress:port",
  "integration_key": "string",
```

```
  "is_enabled": 1,
  "max_prompts": 1,
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "push_info": 1,
  "secret_key": "string",
  "status": "OK"
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
```

```
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

SVM name and UUID for which the Duo profile is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#duo]
[.api-collapsible-fifth-title]
duo

Duo profile for the SVM or cluster-management server (Cserver).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|api_host
|string
a|The URL at which the Duo API is hosted.
```

```
|auto_push
|boolean
a|Automatically sends a push notification for authentication when using
Duo.


|comment
|string
a|Comment for the Duo profile.


|fail_mode
|string
a|Determines the behavior of the system when it cannot communicate with
the Duo service.


|fingerprint
|string
a|The SHA fingerprint corresponding to the Duo secret key.


|http_proxy
|string
a|Specifies the HTTP proxy server to be used when connecting to the Duo
service.


|integration_key
|string
a|The Integration Key associated with the Duo profile.


|is_enabled
|boolean
a|Indicates whether the Duo authentication feature is active or inactive.


|max_prompts
|integer
a|The maximum number of authentication attempts allowed for a user before
the process is terminated.


|owner
|link:#owner[owner]
a|SVM name and UUID for which the Duo profile is configured.
```

```
|push_info
|boolean
a|Additional information sent along with the push notification for Duo
authentication.


|secret_key
|string
a|The Secret Key associated with the Duo profile.


|status
|string
a|Information on the reachability status of Duo.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Update the user account password

:leveloffset: +1


[[IDde8a797a9b5d76a0830ecaf3f5fc1e3d]]
= Update the user account password
```

```
== Overview

This API changes the password for a local user account.

Only cluster administrators with the _"admin"_ role can change the
password for other cluster or SVM user accounts. If you are not a cluster
administrator, you can only change your own password.

== Examples

=== Changing the password of another cluster or SVM user account by a
cluster administrator

Specify the user account name and the new password in the body of the POST
request. The owner.uuid or owner.name are not required to be specified for
a cluster-scoped user account.

For an SVM-scoped account, along with new password and user account name,
specify either the SVM name as the owner.name or SVM uuid as the
owner.uuid in the body of the POST request. These indicate the SVM for
which the user account is created and can be obtained from the response
body of a GET request performed on the _/api/svm/svms_ API.

----

# The API:
POST "/api/security/authentication/password"

# The call to change the password of another cluster user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"name":"cluster_user1","password":"<USER-PASSWORD>"}'
----

----

# The call to change the password of another SVM user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"owner.name":"svm1","name":"svm_user1","password":"<USER-PASSWORD>"}'
----

----

# The call to change the password hash algorithm of the cluster user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"name":"cluster_user1","password":"<USER-
PASSWORD>","password_hash_algorithm":"sha256"}'
```

```
----

----

# The call to change the password hash algorithm of another SVM user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"owner.name":"svm1","name":"svm_user1","password":"<USER-
PASSWORD>","password_hash_algorithm":"sha256"}'
----

=== Changing the password of an SVM-scoped user

NOTE: The IP address in the URI must be same as one of the interfaces
owned by the SVM.

----

# The API:
POST "/api/security/authentication/password"

# The call:
curl -X POST "https://<SVM-ip>/api/security/authentication/password" -d
'{"name":"svm_user1","password":"<USER-PASSWORD>"}'
----

'''




[[ID503241a0c6eaf625dd793e007c1a420d]]
= Update the user account password

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/password`#

*Introduced In:* 9.6

Updates the password for a user account.

== Required parameters

* `name` - User account name.
* `password` - New password for the user account.

== Optional parameters
```

```
* `owner.name` or `owner.uuid` - Name or UUID of the SVM for an SVM-scoped
user account.
* `password_hash_algorithm` - Optional property that specifies the
password hash algorithm used to generate a hash of the user's password for
password matching. Default value is "sha512".

== Related ONTAP commands

* `security login password`

== Learn more

*
xref:{relative_path}security_authentication_password_endpoint_overview.htm
l[DOC /security/authentication/password]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|The user account name whose password is being modified.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the user account. This
field is optional and valid only when a cluster administrator is executing
the API to uniquely identify the account whose password is being modified.
The "owner" field is not required to be specified for SVM user accounts
trying to modify their password.


|password
|string
a|The password string


|password_hash_algorithm
|string
a|Optional property that specifies the password hash algorithm used to
generate a hash of the user's password for password matching.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "password": "string",
  "password_hash_algorithm": "sha512"
}
====
```

```
== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 7077918
| The password cannot contain the username.

| 7077919
| The minimum length for new password does not meet the policy.

| 7077920
| The new password must have both letters and numbers.

| 7077921
| The minimum number of special characters required do not meet the
policy.

| 7077924
```

```
| The new password must be different than last N passwords.

| 7077925
| The new password must be different to the old password.

| 7077940
| The password exceeds maximum supported length.

| 7077941
| Defined password composition exceeds the maximum password length of 128
characters.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
```

```
owner

Owner name and UUID that uniquely identifies the user account. This field
is optional and valid only when a cluster administrator is executing the
API to uniquely identify the account whose password is being modified. The
"owner" field is not required to be specified for SVM user accounts trying
to modify their password.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#account_password]
[.api-collapsible-fifth-title]
account_password

The password object


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The user account name whose password is being modified.
```

```
|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the user account. This
field is optional and valid only when a cluster administrator is executing
the API to uniquely identify the account whose password is being modified.
The "owner" field is not required to be specified for SVM user accounts
trying to modify their password.


|password
|string
a|The password string


|password_hash_algorithm
|string
a|Optional property that specifies the password hash algorithm used to
generate a hash of the user's password for password matching.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage authentication keys (for admins)

:leveloffset: +1


[[IDd082ccd117f88af30606323cb1fb0ce6]]
= Manage authentication keys (for admins)
```

```
== Overview

This API configures the public keys for user accounts.

For secure shell (SSH) access, public-private key pair based
authentication is possible by associating the public key with a user
account.
Prerequisites:
You must have generated the SSH key.
You must be a cluster or SVM administrator to perform the user's public
key.

== Examples

=== Creating a public key for cluster-scoped user accounts

Specify the user account name, public key, index, comment, and optionally
the certificate in the body of the POST request. The owner.uuid or
owner.name are not required for a cluster-scoped user account.

----

# The API:
POST "/api/security/authentication/publickey"

# The call
curl -k https://<mgmt-ip>/api/security/authentication/publickeys --request
POST --data '{ "account": "pubuser2","comment": "Cserver-
Creation","index": 0, "certificate": "<CERTIFICATE-CONTENT>",
"public_key": "<SSH-PUBLIC-KEY>" }'
----

=== Creating a public key for SVM-scoped user accounts

For a SVM-scoped account, specify either the SVM name as the owner.name or
the SVM UUID as the owner.uuid along with other parameters for the user
account. These parameters indicate the SVM that contains the user account
for the public key being created and can be obtained from the response
body of the GET request performed on the API"/api/svm/svms".

----

# The API:
POST "/api/security/authentication/publickey"
```

```
# The call
curl -k https://<mgmt-ip>/api/security/authentication/publickeys --request
POST --data '{ "account": "pubuser4","comment": "Vserver-
Creation","index": 0, "certificate": "<CERTIFICATE-CONTENT>",
"owner.uuid":"513a78c7-8c13-11e9-8f78-
005056bbf6ac","owner.name":"vs0","public_key": "<SSH-PUBLIC-KEY>" }'
----

=== Retrieving the configured public key for user accounts

Retrieves all public keys associated with the user accounts or a filtered
list (for a specific user account name, a specific SVM and so on) of
public keys.

----

# The API:
GET "/api/security/authentication/publickeys"

# The call to retrieve all the user accounts configured in the cluster:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys
----




[[IDefaecfc56307cdcec9ca9c279865320b]]
= Retrieve public keys configured for user accounts

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/publickeys`#

*Introduced In:* 9.7

Retrieves the public keys configured for user accounts.

== Related ONTAP commands

* `security login publickey show`

== Learn more

*
xref:{relative_path}security_authentication_publickeys_endpoint_overview.h
tml[DOC /security/authentication/publickeys]
```

```
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|account.name
|string
|query
|False
a|Filter by account.name


|certificate_expired
|string
|query
|False
a|Filter by certificate_expired

* Introduced in: 9.13


|public_key
|string
|query
|False
a|Filter by public_key


|obfuscated_fingerprint
|string
|query
|False
a|Filter by obfuscated_fingerprint


|certificate_details
|string
```

```
|query
|False
a|Filter by certificate_details

* Introduced in: 9.13


|certificate_revoked
|string
|query
|False
a|Filter by certificate_revoked

* Introduced in: 9.13


|certificate
|string
|query
|False
a|Filter by certificate

* Introduced in: 9.13


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|comment
|string
|query
|False
a|Filter by comment


|sha_fingerprint
```

```
|string
|query
|False
a|Filter by sha_fingerprint


|index
|integer
|query
|False
a|Filter by index

* Max value: 99
* Min value: 0


|scope
|string
|query
|False
a|Filter by scope


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#publickey[publickey]]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "account": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "joe.smith"
      },
      "certificate": "string",
      "certificate_details": "string",
      "certificate_expired": "string",
      "certificate_revoked": "string",
      "comment": "string",
      "obfuscated_fingerprint": "string",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
      },
      "public_key": "string",
      "scope": "string",
      "sha_fingerprint": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the public key.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#publickey]
[.api-collapsible-fifth-title]
publickey

The public key for the user account (to access SSH).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|account
|link:#account_reference[account_reference]
a|

|certificate
|string
a|Optional certificate for the public key.


|certificate_details
```

```
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|index
|integer
a|Index number for the public key (where there are multiple keys for the
same account).


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the public key.


|public_key
|string
a|The public key


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
```

```
|string
a|The SHA fingerprint for the public key (READONLY).


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID3e24e955974a257a91061f8dfe18325c]]
= Create a public key for a user account

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/authentication/publickeys`#

*Introduced In:* 9.7

Creates a public key along with an optional certificate for a user
account.

== Required properties

* `owner.uuid` - UUID of the account owner.
* `name` - User account name.
* `index` - Index number for the public key (where there are multiple keys
for the same account).
* `public_key` - The publickey details for the creation of the user
account.

== Optional properties

* `comment` - Comment text for the public key.
* `certificate` - The certificate in PEM format.

== Related ONTAP commands
```

```
* `security login publickey create`

== Learn more

*
xref:{relative_path}security_authentication_publickeys_endpoint_overview.h
tml[DOC /security/authentication/publickeys]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|
```

```
|certificate
|string
a|Optional certificate for the public key.


|certificate_details
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|index
|integer
a|Index number for the public key (where there are multiple keys for the
same account).


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the public key.


|public_key
|string
a|The public key


|scope
```

```
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|The SHA fingerprint for the public key (READONLY).


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "account": {
    "name": "joe.smith"
  },
  "certificate": "string",
  "certificate_details": "string",
  "certificate_expired": "string",
  "certificate_revoked": "string",
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
====

== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5832705
| Public key already exists for the given user and application.

| 5832707
| Failed to generate fingerprint for the public key.

| 5832722
| The public key cannot be associated with this user on the SVM because a
login method using the given application and authentication method does
not exist for this user.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner
```

```
Owner name and UUID that uniquely identifies the public key.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#publickey]
[.api-collapsible-fifth-title]
publickey

The public key for the user account (to access SSH).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|

|certificate
|string
a|Optional certificate for the public key.


|certificate_details
```

```
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|index
|integer
a|Index number for the public key (where there are multiple keys for the
same account).


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the public key.


|public_key
|string
a|The public key


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
```

```
|string
a|The SHA fingerprint for the public key (READONLY).


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage authentication keys (end-users)

:leveloffset: +1


[[IDfafc22fae4b405ad160d83fb16f33dcc]]
= Manage authentication keys (end-users)



== Overview

This API configures the public keys for end-user (non-cluster admin)
accounts.
Specify the owner UUID, the user account name, and the index in the URI
path. The owner UUID corresponds to the UUID of the SVM containing the
user account associated with the public key and can be obtained from the
response body of the GET request performed on the API "`/api/svm/svms".

The index value corresponds to the public key that needs to be modified or
deleted (it is possible to create more than one public key for the same
user account).
```

```
== Examples

=== Retrieving the specific configured public key for user accounts

----

# The API:
GET
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{inde
x}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/513a78c7-
8c13-11e9-8f78-005056bbf6ac/pubuser4/0
----

=== Updating the public key, certificate, and comment for user accounts

----

# The API:
PATCH
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{inde
x}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/d49de271-
8c11-11e9-8f78-005056bbf6ac/pubuser1/0 --request PATCH --data '{
"comment": "Cserver-modification", "certificate": "<CERTIFICATE-CONTENT>"
, "public_key": "<SSH-PUBLIC-KEY>" }'
----

=== Deleting the public key for user accounts

----

# The API:
DELETE
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{inde
x}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/d49de271-
8c11-11e9-8f78-005056bbf6ac/pubuser1/0 --request DELETE
----
```

```
[[ID6dad5acf3982727ec5a885637c7875fb]]
= Delete a public key for a user account

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/authentication/publickeys/{owner.uuid}/{account.name}/{i
ndex}`#

*Introduced In:* 9.7

Deletes the public key for a user account.

== Related ONTAP commands

* `security login publickey delete`

== Learn more

*
xref:{relative_path}security_authentication_publickeys_owner.uuid_account.
name_index_endpoint_overview.html[DOC
/security/authentication/publickeys/{owner.uuid}/{account.name}/\{index}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID


|account.name
```

```
|string
|path
|True
a|User account name


|index
|integer
|path
|True
a|Index number for the public key (where there are multiple keys for the
same account).

* Max value: 99
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDdc087b11cde74a04683d9149093bace8]]
= Retrieve public keys configured for a user account

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/authentication/publickeys/{owner.uuid}/{account.name}/{i
ndex}`#

*Introduced In:* 9.7

Retrieves the public keys configured for a user account.

== Related ONTAP commands
```

```
* `security login publickey show`

== Learn more

*
xref:{relative_path}security_authentication_publickeys_owner.uuid_account.
name_index_endpoint_overview.html[DOC
/security/authentication/publickeys/{owner.uuid}/{account.name}/\{index}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID


|account.name
|string
|path
|True
a|User account name


|index
|integer
|path
|True
a|Index number for the public key (where there are multiple keys for the
same account).

* Max value: 99
* Min value: 0
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|account
|link:#account_reference[account_reference]
a|

|certificate
|string
a|Optional certificate for the public key.


|certificate_details
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|index
|integer
```

```
a|Index number for the public key (where there are multiple keys for the
same account).


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the public key.


|public_key
|string
a|The public key


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|The SHA fingerprint for the public key (READONLY).


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "account": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
      }
    },
    "name": "joe.smith"
  },
  "certificate": "string",
  "certificate_details": "string",
  "certificate_expired": "string",
  "certificate_revoked": "string",
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the public key.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID44f9cc7ab0fffd3e455d94e27913a35]]
= Update a public key for a user account

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/authentication/publickeys/{owner.uuid}/{account.name}/{i
ndex}`#

*Introduced In:* 9.7

Updates the public key and/or certificate for a user account.

== Related ONTAP commands

* `security login publickey modify`

== Learn more
```

```
*
xref:{relative_path}security_authentication_publickeys_owner.uuid_account.
name_index_endpoint_overview.html[DOC
/security/authentication/publickeys/{owner.uuid}/{account.name}/\{index}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID


|account.name
|string
|path
|True
a|User account name


|index
|integer
|path
|True
a|Index number for the public key (where there are multiple keys for the
same account).

* Max value: 99
* Min value: 0

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|string
a|Optional certificate for the public key.


|certificate_details
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|public_key
|string
a|The public key


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
```

```
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|The SHA fingerprint for the public key (READONLY).


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": "string",
  "certificate_details": "string",
  "certificate_expired": "string",
  "certificate_revoked": "string",
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5832707
| Failed to generate fingerprint for the public key.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|User account
```

```
|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the public key.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#publickey]
[.api-collapsible-fifth-title]
publickey

The public key for the user account (to access SSH).


[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|string
a|Optional certificate for the public key.
```

```
|certificate_details
|string
a|The details present in the certificate (READONLY).


|certificate_expired
|string
a|The expiration details of the certificate (READONLY).


|certificate_revoked
|string
a|The revocation details of the certificate (READONLY).


|comment
|string
a|Optional comment for the public key.


|obfuscated_fingerprint
|string
a|The obfuscated fingerprint for the public key (READONLY).


|public_key
|string
a|The public key


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|The SHA fingerprint for the public key (READONLY).


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage AWS KMS

:leveloffset: +1


[[IDef8150b0491e80dc29a07a8654cbdac7]]
= Manage AWS KMS



== Overview

Amazon Web Services Key Management Services (AWS KMS) is a cloud key
management service (KMS) that provides a secure store for secrets. This
feature allows ONTAP to securely store its encryption keys using AWS KMS.
In order to use AWS KMS with ONTAP, you must first create a
Customer Master Key (CMK) in AWS KMS and provide an Access Key ID and
Secret Access Key for a user that has appropriate access to the newly
created CMK in the AWS KMS."

== Examples

=== Enabling AWS KMS for an SVM

The following example shows how to enable AWS KMS at the SVM-scope. Note
the _return_records=true_ query parameter is used to obtain the newly
created key manager configuration.

----

# The API:
POST /api/security/aws-kms
```

```
# The call:
curl -X POST 'https://<mgmt-ip>/api/security/aws-kms?return_records=true'
-H 'accept: application/hal+json' -d '{"svm":{"uuid":"f36ff553-e713-11ea-
bd56-005056bb4222" }, "region": "us-east-1", "key_id": "kmip-aws",
"access_key_id": "<AWS-ACCESS-KEY-ID>", "secret_access_key": "<AWS-SECRET-
ACCESS-KEY>"}'

# The response:
{
"num_records": 1,
"records": [
    {
    "uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
    "svm": {
        "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
        "name": "vs0"
    },
    "region": "us-east-1",
    "key_id": "kmip-aws",
    "access_key_id": "<AWS-ACCESS-KEY-ID>",
    "_links": {
        "self": {
        "href": "/api/security/aws-kms/f72098a2-e908-11ea-bd56-
005056bb4222"
            }
        }
    }
]
}
----

'''

=== Retrieving all AWS KMS configurations

The following example shows how to retrieve all AWS KMS configurations.

----

# The API:
GET /api/security/aws-kms

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/aws-kms?fields=*'
```

```
# The response:
{
"records": [
    {
    "uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
    "scope": "svm",
    "svm": {
        "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
        "name": "vs0"
    },
    "region": "us-east-1",
    "key_id": "kmip-aws",
    "access_key_id": "<AWS-ACCESS-KEY-ID>",
    "service": "KMS",
    "default_domain": "amazonaws.com",
    "polling_period": 60,
    "timeout": 10,
    "_links": {
        "self": {
        "href": "/api/security/aws-kms/f72098a2-e908-11ea-bd56-
005056bb4222"
            }
        }
    }
],
"num_records": 1,
"_links": {
    "self": {
    "href": "/api/security/aws-kms?fields=*"
    }
}
}
----

'''

=== Retrieving a specific AWS KMS configuration

The following example shows how to retrieve information for a specific AWS
KMS configuration.

----

# The API:
GET /api/security/aws-kms/{uuid}
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/security/aws-kms/f72098a2-e908-11ea-
bd56-005056bb4222?fields=*'

# The response:
{
"uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
"scope": "svm",
"svm": {
    "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
    "name": "vs0"
},
"region": "us-east-1",
"key_id": "kmip-aws",
"access_key_id": "<AWS-ACCESS-KEY-ID>",
"service": "KMS",
"default_domain": "amazonaws.com",
"polling_period": 60,
"timeout": 10,
"_links": {
    "self": {
    "href": "/api/security/aws-kms/f72098a2-e908-11ea-bd56-005056bb4222"
    }
}
}
----

'''

=== Retrieving the advanced properties of an AWS configured for a specific
SVM

These values are not retrieved by default with the 'fields=*' option.
The following example retrieves the advanced properties of a configured
AWS for a specific SVM; there is an added computational cost in retrieving
their values. The properties are not populated for either a collection GET
or an instance GET unless they are explicitly requested using the `fields`
query parameter or GET for all advanced properties is enabled.

----

# The API:
GET /api/security/aws-kms

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/aws-kms/7052c6c0-a503-11ec-
```

```
a68f-005056ac75a0/?fields=state,amazon_reachability,ekmip_reachability'

# The response:
{
"uuid": "d70efc34-aa13-11ec-a059-005056ac7c32",
"state": {
  "cluster_state": true,
  "message": "",
  "code": "0"
},
"amazon_reachability": {
  "reachable": true,
  "message": "",
  "code": "0"
},
"ekmip_reachability": [
  {
    "reachable": true,
    "message": "",
    "code": "0",
    "node": {
      "uuid": "817f544f-a98d-11ec-ae20-005056ac7c32",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/817f544f-a98d-11ec-ae20-
005056ac7c32"
        }
      }
    }
  },
  {
    "reachable": true,
    "message": "",
    "code": "0",
    "node": {
      "uuid": "84b3f5f3-a98d-11ec-9ff4-005056acfbfe",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/84b3f5f3-a98d-11ec-9ff4-
005056acfbfe"
        }
      }
    }
  }
```

```
],
"_links": {
  "self": {
    "href": "/api/security/aws-kms/d70efc34-aa13-11ec-a059-005056ac7c32"
  }
}
}
----

'''

=== Updating the "access_key_id" of a specific AWS KMS configuration

The following example shows how to update the "access_key_id" for a
specific AWS KMS configuration.

----

# The API:
PATCH /api/security/aws-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/aws-kms/f72098a2-e908-11ea-
bd56-005056bb4222/' -d '{"access_key_id": "<AWS-ACCESS-KEY-ID>",
"secret_access_key": "<AWS-SECRET-ACCESS-KEY>"}'
----

'''

=== Updating a specific AWS KMS configuration to allow it to use a proxy.

The following example shows how to update a specific AWS KMS configuration
to allow the AWS KMS instance to use a proxy.

----

# The API:
PATCH /api/security/aws-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/aws-kms/f72098a2-e908-11ea-
bd56-005056bb4222/' -d '{"default_domain": "216.9", "host":
"172.20.216.9", "port": 8000, "service": "10", "verify_host": false,
"verify_ip": false}'
----

'''
```

```
=== Deleting a specific AWS KMS configuration

The following example shows how to delete a specific AWS KMS
configuration.

----

# The API:
DELETE /api/security/aws-kms/{uuid}

# The call;
curl -X DELETE 'https://<mgmt-ip>/api/security/aws-kms/f72098a2-e908-11ea-
bd56-005056bb4222'
----

'''

=== Restoring keys from a KMIP server

The following example shows how to restore keys for a AWS KMS
configuration.

----

# The API:
POST /api/security/aws-kms/{uuid}/restore

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/aws-kms/33820b57-ec90-11ea-
875e-005056bbf3f0/restore'
----

'''





[[ID1067638eb9db9c21f2c28d6ef70191e6]]
= Retrieve all AWS KMS instances configured for all clusters and SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/aws-kms`#

*Introduced In:* 9.12
```

```
Retrieves all AWS KMS instances configured for all clusters and SVMs.

== Related ONTAP commands

* `security key-manager external aws show`
* `security key-manager external aws check`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|state.message
|string
|query
|False
a|Filter by state.message


|state.cluster_state
|boolean
|query
|False
a|Filter by state.cluster_state


|state.code
|string
|query
|False
a|Filter by state.code


|polling_period
|integer
|query
|False
a|Filter by polling_period
```

```
|proxy_port
|integer
|query
|False
a|Filter by proxy_port


|verify_ip
|boolean
|query
|False
a|Filter by verify_ip


|host
|string
|query
|False
a|Filter by host


|proxy_type
|string
|query
|False
a|Filter by proxy_type


|port
|integer
|query
|False
a|Filter by port


|timeout
|integer
|query
|False
a|Filter by timeout


|uuid
|string
|query
|False
a|Filter by uuid
```

```
|service
|string
|query
|False
a|Filter by service


|ekmip_reachability.message
|string
|query
|False
a|Filter by ekmip_reachability.message


|ekmip_reachability.node.name
|string
|query
|False
a|Filter by ekmip_reachability.node.name


|ekmip_reachability.node.uuid
|string
|query
|False
a|Filter by ekmip_reachability.node.uuid


|ekmip_reachability.reachable
|boolean
|query
|False
a|Filter by ekmip_reachability.reachable


|ekmip_reachability.code
|string
|query
|False
a|Filter by ekmip_reachability.code


|region
|string
|query
```

```
|False
a|Filter by region


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|verify_host
|boolean
|query
|False
a|Filter by verify_host


|encryption_context
|string
|query
|False
a|Filter by encryption_context


|proxy_username
|string
|query
|False
a|Filter by proxy_username


|amazon_reachability.message
|string
|query
|False
a|Filter by amazon_reachability.message


|amazon_reachability.code
```

```
|string
|query
|False
a|Filter by amazon_reachability.code


|amazon_reachability.reachable
|boolean
|query
|False
a|Filter by amazon_reachability.reachable


|default_domain
|string
|query
|False
a|Filter by default_domain


|access_key_id
|string
|query
|False
a|Filter by access_key_id


|verify
|boolean
|query
|False
a|Filter by verify


|scope
|string
|query
|False
a|Filter by scope


|proxy_host
|string
|query
|False
a|Filter by proxy_host
```

```
|skip_verify
|boolean
|query
|False
a|Filter by skip_verify


|key_id
|string
|query
|False
a|Filter by key_id


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
```

```
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#aws_kms[aws_kms]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_key_id": "<AWS-ACCESS-KEY-ID>",
      "amazon_reachability": {
        "code": "346758",
        "message": "Amazon KMS is not reachable from all nodes -
<reason>."
      },
      "default_domain": "domainName",
      "ekmip_reachability": [
        {
          "code": "346758",
          "message": "embedded KMIP server status unavailable on node.",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "encryption_context": "aws:fsx:fs-id=fs-0785c8beceb895999",
      "host": "aws-host.host.com",
      "key_id": "kmip-aws",
      "polling_period": 55,
      "port": 443,
      "proxy_host": "proxy.eng.com",
      "proxy_port": 1234,
      "proxy_type": "http",
```

```
      "proxy_username": "proxyuser",
      "region": "us-east-1",
      "scope": "string",
      "service": "dynamodb.*.amazonaws.com",
      "skip_verify": "",
      "state": {
        "code": "346758",
        "message": "AWS KMS key protection is unavailable on the following
nodes: node1, node2."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "timeout": 20,
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "verify": "",
      "verify_host": 1,
      "verify_ip": ""
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.
```

```
| 65537915
| The Amazon Web Service Key Management Service is unreachable from one or
more nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#amazon_reachability]
[.api-collapsible-fifth-title]
amazon_reachability

Indicates whether or not the Amazon KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Amazon KMS is
reachable from all nodes in the cluster.


|message
|string
a|Error message returned when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Amazon KMS is reachable from all nodes of the
cluster.
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when AWS KMS key protection is available on all nodes of the
cluster.


|code
```

```
|string
a|Code corresponding to the message. Returns a 0 if AWS KMS key protection
is available on all nodes of the cluster.


|message
|string
a|Error message set when cluster_state is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#aws_kms]
[.api-collapsible-fifth-title]
aws_kms
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.


|key_id
```

```
|string
a|AWS Key ID.


|polling_period
|integer
a|Polling period in minutes.


|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.


|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|service
|string
```

```
a|AWS service type.


|skip_verify
|boolean
a|Set to true to bypass verification of the user provided access_key_id
and secret_access_key. An error will be returned if 'skip_verify' is
provided but 'access_key_id' is not.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|AWS Connection timeout, in seconds.


|uuid
|string
a|A unique identifier for the AWS KMS.


|verify
|boolean
a|Set to true to verify the AWS KMS host.


|verify_host
|boolean
a|Set to true to verify the AWS KMS host's hostname.


|verify_ip
|boolean
a|Set to true to verify the AWS KMS host's IP address.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDeae15bb3156be563d6b1e319d5bc3711]]
= Configure the AWS KMS configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/aws-kms`#

*Introduced In:* 9.12

Configures the AWS KMS configuration for the specified SVM.

== Required properties

* `access_key_id` - AWS access key ID of the user who has the appropriate
access to AWS KMS.
* `secret_access_key` - AWS secret access key for the access key ID
provided.
* `svm.uuid` or `svm.name` - Existing SVM in which to create an AWS KMS.
* `region` - AWS region of the AWS KMS.
* `key_id` - AWS Key ID

== Optional properties

* `service` - AWS service type.
* `default_domain` - AWS KMS default domain.
* `host` - AWS KMS host's hostname.
* `port` - AWS KMS port.
* `proxy_type` - Type of proxy (http, https, etc.), if proxy configuration
is used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
```

```
* `proxy_password` - Proxy password if proxy configuration is used.
* `polling_period` - Polling period in minutes.
* `encryption_context` - Additional layer of authentication and logging.

== Related ONTAP commands

* `security key-manager external aws enable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.
```

```
|key_id
|string
a|AWS Key ID.


|polling_period
|integer
a|Polling period in minutes.


|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.


|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.
```

```
|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|secret_access_key
|string
a|AWS Secret Access Key for the provided access key ID.


|service
|string
a|AWS service type.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the AWS KMS.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access_key_id": "<AWS-ACCESS-KEY-ID>",
  "amazon_reachability": {
    "code": "346758",
    "message": "Amazon KMS is not reachable from all nodes - <reason>."
  },
  "default_domain": "domainName",
  "ekmip_reachability": [
```

```
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "encryption_context": "aws:fsx:fs-id=fs-0785c8beceb895999",
  "host": "aws-host.host.com",
  "key_id": "kmip-aws",
  "polling_period": 55,
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_password": "awskze-Jwjje2-WJJPer",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "region": "us-east-1",
  "scope": "string",
  "secret_access_key": "<AWS-SECRET-ACCESS-KEY>",
  "service": "dynamodb.*.amazonaws.com",
  "skip_verify": "",
  "state": {
    "code": "346758",
    "message": "AWS KMS key protection is unavailable on the following
nodes: node1, node2."
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "verify": "",
  "verify_host": 1,
  "verify_ip": ""
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 3735622
| Certificate type not supported for create operation.

| 3735645
| You cannot specify a value for serial as it is generated automatically.

| 3735657
| Specifying \"-subtype\" when creating a certificate is not supported.

| 3735664
| Specified key size is not supported in FIPS mode.

| 3735665
| Specified hash function is not supported in FIPS mode.

| 3735700
| Specified key size is not supported.

| 65536600
| Nodes are out of quorum.

| 65537518
| Failed to find a LIF with Cluster role on node. One or more nodes may be
out of quorum.

| 65537900
| Failed to enable the Amazon Web Service Key Management Service for an
SVM due to an invalid secret access key.

| 65537901
| The Amazon Web Service Key Management Service (AWSKMS) cannot be enabled
because all nodes in the cluster are not running a version that supports
the AWSKMS feature.

| 65537906
| Failed to store the secret access key.
```

```
| 65537907
| The Amazon Web Service Key Management Service is disabled on the
cluster. For further assistance, contact technical support.

| 65537908
| The Amazon Web Service Key Management Service is not supported for the
admin SVM.

| 65537910
| Failed to configure Amazon Web Service Key Management Service for an SVM
because a key manager has already been configured for the SVM.

| 65537911
| The Amazon Web Service Key Management Service is not supported in
MetroCluster configurations.

| 65537912
| The Amazon Web Service Key Management Service cannot be configured for
an SVM because one or more volume encryption keys of the SVM are stored on
the admin SVM.

| 65537926
| The Amazon Web Service Key Management Service is not configured for this
SVM.

| 65539704
| The key manager cannot be configured because the SVM has NAE volumes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#amazon_reachability]
[.api-collapsible-fifth-title]
amazon_reachability

Indicates whether or not the Amazon KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Amazon KMS is
reachable from all nodes in the cluster.


|message
|string
a|Error message returned when 'reachable' is false.


|reachable
|boolean
```

```
a|Set to true if the Amazon KMS is reachable from all nodes of the
cluster.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when AWS KMS key protection is available on all nodes of the
cluster.


|code
```

```
|string
a|Code corresponding to the message. Returns a 0 if AWS KMS key protection
is available on all nodes of the cluster.


|message
|string
a|Error message set when cluster_state is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#aws_kms]
[.api-collapsible-fifth-title]
aws_kms

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.


|key_id
|string
a|AWS Key ID.
```

```
|polling_period
|integer
a|Polling period in minutes.


|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.


|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|secret_access_key
```

```
|string
a|AWS Secret Access Key for the provided access key ID.


|service
|string
a|AWS service type.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the AWS KMS.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6bbfa50420ed4c5e76ac7249528c79d2]]
= Re-key or re-version an AWS KMS key encryption key for AWS KMS

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/aws-kms/{aws_kms.uuid}/rekey-external`#

*Introduced In:* 9.12

Rekeys or re-versions the AWS KMS Key Encryption Key (KEK) for the given
AWS KMS.

== Related ONTAP commands

* `security key-manager external aws rekey-external`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aws_kms.uuid
|string
|path
|True
a|UUID of the existing AWS KMS configuration.
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier of the AWS KMS key encryption key.


|===


.Example request
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "key_id": "key01"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537538
| Internal error. Failed to get unwrapped key for a given key ID.

| 65537543
| Internal Error. Missing top-level internal key protection key (KEK) on a
node.

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager before running the rekey operation.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65537919
| External rekey failed on one or more nodes.

| 65537926
| AWS KMS is not configured for the given SVM.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#aws_kms_key]
[.api-collapsible-fifth-title]
aws_kms_key

[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier of the AWS KMS key encryption key.


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6463420411e519aa22f53e3c45354a44]]
= Re-key SVM KEK for an AWS KMS

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/aws-kms/{aws_kms.uuid}/rekey-internal`#

*Introduced In:* 9.12

Rekeys SVM KEK for the given AWS KMS.

== Related ONTAP commands

* `security key-manager external aws rekey-internal`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aws_kms.uuid
|string
|path
|True
a|UUID of the existing AWS KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537556
| Unable to successfully encrypt or decrypt because the configured
external key manager for the given SVM is in a blocked state.

| 65537559
| There are no existing internal keys for the SVM. A rekey operation is
allowed for an SVM with one or more encryption keys.

| 65537566
| Internal error. All nodes in the cluster are not currently online.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65537926
| AWS KMS is not configured for the given SVM.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
```

```
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
```

```
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID7078eeeab51505976612f3b37f780058]]
= Restore keys for an SVM from a configured AWS KMS

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/aws-kms/{aws_kms.uuid}/restore`#

*Introduced In:* 9.12

Restores the keys for an SVM from a configured AWS KMS.

== Related ONTAP commands

* `security key-manager external AWS restore`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aws_kms.uuid
```

```
|string
|path
|True
a|UUID of the existing AWS KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536082
| Unable to restore all keys.

| 65537544
| Missing wrapped top-level internal key protection key (KEK) from
internal database.

| 65537926
| The Amazon Web Service Key Management Service is not configured for the
given SVM.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4913d5049bed58d3f85a07d3bcdab68c]]
= Delete an AWS KMS configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/aws-kms/{uuid}`#
```

```
*Introduced In:* 9.12

Deletes an AWS KMS configuration.

== Related ONTAP commands

* `security key-manager external aws disable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|AWS KMS UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536817
| Internal error. Failed to determine if it is safe to disable key
manager.

| 65536827
| Internal error. Failed to determine if the given SVM has any encrypted
volumes.

| 65536834
| Internal error. Failed to get existing key-server details for the given
SVM.

| 65536883
| Internal error. Volume encryption key is missing for a volume.

| 65536884
| Internal error. Volume encryption key is invalid for a volume.

| 65537106
| Volume encryption keys (VEK) for one or more encrypted volumes are
stored on the key manager configured for the given SVM.

| 65537926
| Amazon Web Service Key Management Service is not configured for SVM.

| 196608080
| One or more nodes in the cluster have the root volume encrypted using
NVE (NetApp Volume Encryption).
```

```
| 196608301
| Internal error. Failed to get encryption type.

| 196608332
| NAE aggregates found in the cluster.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4ab91bb1f2cd4465a09b11da59e4f6e9]]
= Retrieve an AWS KMS configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/aws-kms/{uuid}`#

*Introduced In:* 9.12

Retrieves the AWS KMS configuration for the SVM specified by the UUID.

== Related ONTAP commands

* `security key-manager external aws show`
* `security key-manager external aws check`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
```

```
|string
|path
|True
a|AWS KMS UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
```

```
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.


|key_id
|string
a|AWS Key ID.


|polling_period
|integer
a|Polling period in minutes.


|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.
```

```
|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|service
|string
a|AWS service type.


|skip_verify
|boolean
a|Set to true to bypass verification of the user provided access_key_id
and secret_access_key. An error will be returned if 'skip_verify' is
provided but 'access_key_id' is not.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|AWS Connection timeout, in seconds.


|uuid
|string
a|A unique identifier for the AWS KMS.
```

```
|verify
|boolean
a|Set to true to verify the AWS KMS host.


|verify_host
|boolean
a|Set to true to verify the AWS KMS host's hostname.


|verify_ip
|boolean
a|Set to true to verify the AWS KMS host's IP address.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_key_id": "<AWS-ACCESS-KEY-ID>",
  "amazon_reachability": {
    "code": "346758",
    "message": "Amazon KMS is not reachable from all nodes - <reason>."
  },
  "default_domain": "domainName",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "encryption_context": "aws:fsx:fs-id=fs-0785c8beceb895999",
  "host": "aws-host.host.com",
  "key_id": "kmip-aws",
  "polling_period": 55,
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "region": "us-east-1",
  "scope": "string",
  "service": "dynamodb.*.amazonaws.com",
  "skip_verify": "",
  "state": {
    "code": "346758",
    "message": "AWS KMS key protection is unavailable on the following
nodes: node1, node2."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timeout": 20,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "verify": "",
  "verify_host": 1,
  "verify_ip": ""
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.

| 65537915
| The Amazon Web Service Key Management Service is unreachable from one or
more nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
```

```
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#amazon_reachability]
[.api-collapsible-fifth-title]
amazon_reachability

Indicates whether or not the Amazon KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Amazon KMS is
reachable from all nodes in the cluster.


|message
|string
a|Error message returned when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Amazon KMS is reachable from all nodes of the
cluster.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
```

```
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when AWS KMS key protection is available on all nodes of the
cluster.


|code
|string
a|Code corresponding to the message. Returns a 0 if AWS KMS key protection
is available on all nodes of the cluster.


|message
|string
a|Error message set when cluster_state is false.
```

```
|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID15b862beb40e4e5fdd65f03f0fe33aa9]]
= Update an AWS KMS configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/aws-kms/{uuid}`#

*Introduced In:* 9.12

Updates the AWS KMS configuration.

== Optional properties

* `region` - AWS region of the AWS KMS.
* `service` - AWS service type.
* `default_domain` - AWS KMS default domain.
* `port` - AWS KMS port.
* `proxy_type` - Type of proxy (http, https, etc.), if proxy configuration
is used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `polling_period` - Polling period in minutes.
* `timeout` - AWS Connection timeout, in seconds.
* `verify` - Set to true to verify the AWS KMS host.
* `verify_host` - Set to true to verify the AWS KMS host's hostname.
* `verify_ip` - Set to true to verify the AWS KMS host's IP address.
* `host` - AWS KMS host's hostname.
* `secret_access_key` - AWS secret access key for the access key ID
provided.
* `access_key_id` - AWS access key ID of the user with the appropriate
access to AWS KMS.
* `skip_verify` - Set to true to bypass verification of the user provided
access_key_id and secret_access_key.
* `encryption_context` - Additional layer of authentication and logging.

== Related ONTAP commands

* `security key-manager external aws update-config`
* `security key-manager external aws update-credentials`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|AWS KMS UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
```

```
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.


|polling_period
|integer
a|Polling period in minutes.


|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.
```

```
|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.


|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|secret_access_key
|string
a|AWS Secret Access Key for the provided access key ID.


|service
|string
a|AWS service type.


|skip_verify
|boolean
a|Set to true to bypass verification of the user provided access_key_id
and secret_access_key. An error will be returned if 'skip_verify' is
```

```
provided but 'access_key_id' is not.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|timeout
|integer
a|AWS Connection timeout, in seconds.


|uuid
|string
a|A unique identifier for the AWS KMS.


|verify
|boolean
a|Set to true to verify the AWS KMS host.


|verify_host
|boolean
a|Set to true to verify the AWS KMS host's hostname.


|verify_ip
|boolean
a|Set to true to verify the AWS KMS host's IP address.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access_key_id": "<AWS-ACCESS-KEY-ID>",
  "amazon_reachability": {
    "code": "346758",
    "message": "Amazon KMS is not reachable from all nodes - <reason>."
  },
```

```
  "default_domain": "domainName",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "encryption_context": "aws:fsx:fs-id=fs-0785c8beceb895999",
  "host": "aws-host.host.com",
  "polling_period": 55,
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_password": "awskze-Jwjje2-WJJPer",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "region": "us-east-1",
  "scope": "string",
  "secret_access_key": "<AWS-SECRET-ACCESS-KEY>",
  "service": "dynamodb.*.amazonaws.com",
  "skip_verify": "",
  "state": {
    "code": "346758",
    "message": "AWS KMS key protection is unavailable on the following
nodes: node1, node2."
  },
  "timeout": 20,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "verify": "",
  "verify_host": 1,
  "verify_ip": ""
}
====

== Response
```


Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537541
| No inputs provided for the REST API PATCH request.

| 65537906
| Failed to store the secret access key.

| 65537920
| Secret access key cannot be empty.

| 65537921
| Unable to connect to the Amazon Web Service Key Management Service
(AWSKMS) using these credentials.

| 65537924
| Access key ID cannot be empty.

| 65537926
| Amazon Web Service Key Management Service is not configured for SVM.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#amazon_reachability]
[.api-collapsible-fifth-title]
amazon_reachability

Indicates whether or not the Amazon KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Amazon KMS is
reachable from all nodes in the cluster.


|message
|string
a|Error message returned when 'reachable' is false.


|reachable
```

```
|boolean
a|Set to true if the Amazon KMS is reachable from all nodes of the
cluster.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when AWS KMS key protection is available on all nodes of the
cluster.
```

```
|code
|string
a|Code corresponding to the message. Returns a 0 if AWS KMS key protection
is available on all nodes of the cluster.


|message
|string
a|Error message set when cluster_state is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#aws_kms]
[.api-collapsible-fifth-title]
aws_kms

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access_key_id
|string
a|AWS Access Key ID of the user that has appropriate access to AWS KMS.


|amazon_reachability
|link:#amazon_reachability[amazon_reachability]
a|Indicates whether or not the Amazon KMS is reachable from all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|default_domain
|string
a|AWS KMS default domain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|encryption_context
|string
a|Additional layer of authentication and logging.


|host
|string
a|AWS KMS host's hostname.


|polling_period
|integer
a|Polling period in minutes.
```

```
|port
|integer
a|AWS KMS port.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Proxy type.


|proxy_username
|string
a|Proxy username.


|region
|string
a|AWS region of the AWS KMS.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|secret_access_key
|string
a|AWS Secret Access Key for the provided access key ID.
```

```
|service
|string
a|AWS service type.


|skip_verify
|boolean
a|Set to true to bypass verification of the user provided access_key_id
and secret_access_key. An error will be returned if 'skip_verify' is
provided but 'access_key_id' is not.


|state
|link:#state[state]
a|Indicates whether or not the Amazon Web Services Key Management Service
(AWS KMS) key protection is available cluster-wide.


|timeout
|integer
a|AWS Connection timeout, in seconds.


|uuid
|string
a|A unique identifier for the AWS KMS.


|verify
|boolean
a|Set to true to verify the AWS KMS host.


|verify_host
|boolean
a|Set to true to verify the AWS KMS host's hostname.


|verify_ip
|boolean
a|Set to true to verify the AWS KMS host's IP address.


|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage Azure Key Vaults

:leveloffset: +1


[[ID5ba9c1ead37fcc5d6fbf924ca62cba04]]
= Manage Azure Key Vaults



== Overview

Azure Key Vault (AKV) is a cloud key management service (KMS) that
provides a secure store for secrets. This feature
allows ONTAP to securely store its encryption keys using AKV.
In order to use AKV with ONTAP, you must first deploy an Azure application
with the appropriate access to an AKV and then provide
ONTAP with the necessary details, such as key vault name, application ID
so that ONTAP can communicate with the deployed Azure application.
The properties "state", "azure_reachability" and "ekmip_reachability" are
considered advanced properties and are populated only when explicitly
requested.

== Examples

=== Enabling an AKV configuration for an SVM using the certificate
authentication method

The example AKV configuration is enabled for a specific SVM. Note the
_return_records=true_ query parameter is used to obtain the newly created
key-manager keystore configuration.
```

```
----

# The API:
POST /api/security/azure-key-vaults

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-
vaults?return_records=true' -H 'accept: application/hal+json' -d "{
\"svm\": { \"uuid\": \"4f7abf4c-9a07-11ea-8d52-005056bbeba5\" },
\"client_id\": \"client1\", \"tenant_id\": \"tenant1\", \"name\":
\"https:://mykeyvault.vault.azure.net/\", \"key_id\": \"https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74\",
\"client_certificate\" : \"<CERTIFICATE-CONTENT>\" }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "024cd3cf-9a08-11ea-8d52-005056bbeba5",
    "svm": {
      "uuid": "4f7abf4c-9a07-11ea-8d52-005056bbeba5",
      "name": "vs0"
    },
    "client_id": "client1",
    "tenant_id": "tenant1",
    "name": "https:://mykeyvault.vault.azure.net/",
    "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
    "_links": {
      "self": {
        "href": "/api/security/azure-key-vaults/024cd3cf-9a08-11ea-8d52-
005056bbeba5"
      }
    }
  }
]
}
----

'''

=== Creating an inactive AKV configuration for an SVM using the client
secret authentication method
```

```
The example AKV configuration is created for a specific SVM but is not
enabled.
Note the _create_inactive=true_ parameter that is used to indicate that
the configuration should be created but not enabled.
Note the _return_records=true_ query parameter is used to obtain the newly
created key-manager keystore configuration.

----

# The API:
POST /api/security/azure-key-vaults

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-
vaults?create_inactive=true&return_records=true' -H 'accept:
application/hal+json' -d "{ \"svm\": { \"uuid\": \"4f7abf4c-9a07-11ea-
8d52-005056bbeba5\" }, \"configuration\": {\"name\": \"myConfiguration\"},
\"client_id\": \"client1\", \"tenant_id\": \"tenant1\", \"name\":
\"https:://mykeyvault.vault.azure.net/\", \"key_id\": \"https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74\",
\"client_secret\" : \"myclientPwd\" }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
    "configuration": {"name": "myConfiguration"},
    "client_id": "client1",
    "tenant_id": "tenant1",
    "name": "https:://mykeyvault.vault.azure.net/",
    "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
    "_links": {
      "self": {
        "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-
005056bbeba5"
      }
    }
  }
]
}
----

'''
```

```
=== Retrieving the AKVs configured for all clusters and SVMs

The following example shows how to retrieve all configured AKVs along with
their configurations.

----

# The API:
GET /api/security/azure-key-vaults

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/azure-key-vaults?fields=*'

# The response:
{
"records": [
  {
    "configuration": {
      "uuid": "024cd3cf-9a08-11ea-8d52-005056bbeba5",
      "name" : "default"
    },
    "uuid": "024cd3cf-9a08-11ea-8d52-005056bbeba5",
    "scope": "svm",
    "svm": {
      "uuid": "4f7abf4c-9a07-11ea-8d52-005056bbeba5",
      "name": "vs0"
    },
    "client_id": "client1",
    "tenant_id": "tenant1",
    "name": "https:://mykeyvault.vault.azure.net/",
    "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
    "authentication_method": "client_secret",
    "enabled": true,
    "_links": {
      "self": {
        "href": "/api/security/azure-key-vaults/024cd3cf-9a08-11ea-8d52-
005056bbeba5"
      }
    }
  },
  {
    "configuration": {
      "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
      "name" : "new-config"
    },
```

```
    "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
    "scope": "cluster",
    "client_id": "client1",
    "tenant_id": "tenant1",
    "name": "https:://mykeyvault.vault.azure.net/",
    "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
    "authentication_method": "certificate",
    "enabled": false,
    "_links": {
      "self": {
        "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-
005056bbeba5"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/azure-key-vaults?fields=*"
  }
}
}
----

'''

=== Retrieving a specific AKV configuration

The following example retrieves a specific AKV configuration.

----

# The API:
GET /api/security/azure-key-vaults

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5?fields=*'

# The response:
{
"configuration": {
  "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
  "name" : "default"
```

```
},
"uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
"scope": "cluster",
"client_id": "client1",
"tenant_id": "tenant1",
"name": "https:://mykeyvault.vault.azure.net/",
"key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
"authentication_method": "client_secret",
"enabled": true,
"_links": {
  "self": {
    "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-
005056bbeba5"
  }
}
}
----

'''

=== Retrieving the advanced properties of a specific, enabled AKV
configuration

The following example retrieves the advanced properties of a specific
enabled AKV configuration (inactive AKV configurations do not have these
advanced properties).

----

# The API:
GET /api/security/azure-key-vaults

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-
005056bbeba5?fields=state,azure_reachability,ekmip_reachability"'
{
"configuration": {
  "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
  "name" : "default"
},
"uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
"name": "https://10.234.237.18",
"state": {
  "cluster_state": true,
```

```
  "message": "",
  "code": "0"
},
"azure_reachability": {
  "reachable": true,
  "message": "",
  "code": "0"
},
"ekmip_reachability": [
  {
    "node": {
        "uuid": "d208115f-7721-11eb-bf83-005056bb150e",
        "name": "node1",
        "_links": {
            "self": {
            "href": "/api/cluster/nodes/d208115f-7721-11eb-bf83-
005056bb150e"
            }
        }
    },
    "reachable": true,
    "message": "",
    "code": "0"
  },
  {
    "node": {
        "uuid": "e208115f-7721-11eb-bf83-005056bb150e",
        "name": "node2",
        "_links": {
            "self": {
            "href": "/api/cluster/nodes/e208115f-7721-11eb-bf83-
005056bb150e"
            }
        }
    },
    "reachable": true,
    "message": "",
    "code": "0"
  }
],
"_links": {
  "self": {
    "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-
005056bbeba5"
  }
}
```

```
}
----

'''

=== Updating the client secret of a specific AKV configuration

The following example updates the client secret of a specific AKV
configuration.

----

# The API:
PATCH /api/security/azure-key-vaults

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5' -d "{ \"client_secret\": \"<NEW-SECRET>\" }"
----

'''

=== Updating the client certificate and key of a specific AKV
configuration

The following example updates the client certificate and key of a
configured AKV for a specific AKV configuration.

----

# The API:
PATCH /api/security/azure-key-vaults

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5' -d "{ \"client_certificate\": \"<CERTIFICATE-
CONTENT>\" }"
----

'''

=== Deleting a specific AKV configuration

The following example deletes a specific, enabled AKV.

----
```

```
# The API:
DELETE /api/security/azure-key-vaults

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5'
----

'''

=== Restoring the keys for a specific AKV configuration

The following example restores all the keys of a specific AKV
configuration.

----

# The API:
POST security/azure-key-vaults/{uuid}/restore

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5/restore'

# The response:
{
"job": {
  "uuid": "6ab6946f-9a0c-11ea-8d52-005056bbeba5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6ab6946f-9a0c-11ea-8d52-005056bbeba5"
    }
  }
}
}
----

'''

=== Rekeying the internal key for a specific AKV configuration

The following example rekeys the internal key of a specific AKV
configuration.

----

# The API:
```

```
POST security/azure-key-vaults/{uuid}/rekey-internal

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-
9a06-11ea-8d52-005056bbeba5/rekey-internal'

# The response:
{
"job": {
  "uuid": "6ab6946f-9a0c-11ea-8d52-005056bbeba5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6ab6946f-9a0c-11ea-8d52-005056bbeba5"
    }
  }
}
}
----

'''





[[ID459aa84538d4a24e45c811d4687dd2e4]]
= Retrieve AKVs configured for all clusters and SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/azure-key-vaults`#

*Introduced In:* 9.8

Retrieves AKVs configured for all clusters and SVMs.

== Related ONTAP commands

* `security key-manager external azure show`
* `security key-manager external azure check`


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|verify_host
|boolean
|query
|False
a|Filter by verify_host

* Introduced in: 9.14


|tenant_id
|string
|query
|False
a|Filter by tenant_id


|oauth_host
|string
|query
|False
a|Filter by oauth_host

* Introduced in: 9.14


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|vault_host
|string
|query
```

```
|False
a|Filter by vault_host

* Introduced in: 9.14


|azure_reachability.reachable
|boolean
|query
|False
a|Filter by azure_reachability.reachable


|azure_reachability.code
|string
|query
|False
a|Filter by azure_reachability.code


|azure_reachability.message
|string
|query
|False
a|Filter by azure_reachability.message


|proxy_host
|string
|query
|False
a|Filter by proxy_host


|scope
|string
|query
|False
a|Filter by scope


|client_id
|string
|query
|False
a|Filter by client_id
```

```
|skip_verification
|boolean
|query
|False
a|Filter by skip_verification

* Introduced in: 9.16


|key_id
|string
|query
|False
a|Filter by key_id


|proxy_username
|string
|query
|False
a|Filter by proxy_username


|proxy_port
|integer
|query
|False
a|Filter by proxy_port


|verify_ip
|boolean
|query
|False
a|Filter by verify_ip

* Introduced in: 9.14


|enabled
|boolean
|query
|False
a|Filter by enabled

* Introduced in: 9.14
```

```
|name
|string
|query
|False
a|Filter by name


|state.code
|string
|query
|False
a|Filter by state.code


|state.message
|string
|query
|False
a|Filter by state.message


|state.available
|boolean
|query
|False
a|Filter by state.available


|configuration.uuid
|string
|query
|False
a|Filter by configuration.uuid

* Introduced in: 9.14


|configuration.name
|string
|query
|False
a|Filter by configuration.name

* Introduced in: 9.14
```

```
|authentication_method
|string
|query
|False
a|Filter by authentication_method

* Introduced in: 9.10


|ekmip_reachability.message
|string
|query
|False
a|Filter by ekmip_reachability.message


|ekmip_reachability.node.name
|string
|query
|False
a|Filter by ekmip_reachability.node.name


|ekmip_reachability.node.uuid
|string
|query
|False
a|Filter by ekmip_reachability.node.uuid


|ekmip_reachability.reachable
|boolean
|query
|False
a|Filter by ekmip_reachability.reachable


|ekmip_reachability.code
|string
|query
|False
a|Filter by ekmip_reachability.code


|uuid
|string
```

```
|query
|False
a|Filter by uuid


|proxy_type
|string
|query
|False
a|Filter by proxy_type


|port
|integer
|query
|False
a|Filter by port

* Introduced in: 9.14


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#azure_key_vault[azure_key_vault]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_method": "client_secret",
      "azure_reachability": {
        "code": "346758",
        "message": "AKV service is not reachable from all nodes - reason."
      },
      "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
      "configuration": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      },
      "ekmip_reachability": [
        {
          "code": "346758",
          "message": "embedded KMIP server status unavailable on node.",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
```

```
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "key_id":
"https://keyvault1.vault.azure.net/keys/key1/12345678901234567890123456789
012",
      "name": "https://kmip-akv-keyvault.vault.azure.net/",
      "oauth_host": "login.microsoftonline.com",
      "port": 443,
      "proxy_host": "proxy.eng.com",
      "proxy_port": 1234,
      "proxy_type": "http",
      "proxy_username": "proxyuser",
      "scope": "string",
      "skip_verification": "",
      "state": {
        "code": "346758",
        "message": "Top-level internal key protection key (KEK) is
unavailable on the following nodes with the associated reasons: Node:
node1. Reason: No volumes created yet for the SVM. Wrapped KEK status will
be available after creating encrypted volumes."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vault_host": "vault.azure.net",
      "verify_host": "",
      "verify_ip": ""
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537549
| The Azure Key Vault Key Management Service is unreachable from one or
more nodes.

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
```

```
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#azure_reachability]
[.api-collapsible-fifth-title]
azure_reachability

Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Code corresponding to the status message. Returns a 0 if AKV service is
reachable from all nodes in the cluster.


|message
|string
a|Error message set when reachability is false.


|reachable
|boolean
a|Set to true when the AKV service is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===
```

```
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the AKV wrapped internal key is available cluster
wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|boolean
a|Set to true when an AKV wrapped internal key is present on all nodes of
```

```
the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV wrapped key
is available on all nodes in the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#azure_key_vault]
[.api-collapsible-fifth-title]
azure_key_vault

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.


|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|
```

```
|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier of AKV key encryption key.


|name
|string
a|Name of the deployed AKV that will be used by ONTAP for storing keys.

* example: https://kmip-akv-keyvault.vault.azure.net/
* format: uri
* Introduced in: 9.8
* readCreate: 1
* x-nullable: true


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Type of proxy.


|proxy_username
```

```
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
a|AKV host subdomain.
```

```
|verify_host
|boolean
a|Verify the identity of the AKV host name.


|verify_ip
|boolean
a|Verify the identity of the AKV IP address.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDaa497e914df6bf6d15c3b4ce88db58a8]]
= Create an AKV configuration for all clusters and SVMs

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/azure-key-vaults`#

*Introduced In:* 9.8

Configures the AKV configuration for all clusters and SVMs.

== Required properties:

* `svm.uuid` or `svm.name` - Existing SVM in which to create a AKV.
* `client_id` - Application (client) ID of the deployed Azure application
with appropriate access to an AKV.
* `tenant_id` - Directory (tenant) ID of the deployed Azure application
with appropriate access to an AKV.
* `client_secret` or `client_certificate` - Secret or PKCS12 Certificate
used by the application to prove its identity to AKV.
```

```
* `key_id`- Key Identifier of AKV encryption key.
* `name` - Name of the deployed AKV used by ONTAP for storing keys.

== Optional properties:

* `port` - Authorization server and vault port number.
* `oauth_host` - Open authorization server host name.
* `vault_host` - AKV host subdomain.
* `proxy_type` - Type of proxy (http, https etc.) if proxy configuration
is used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `configuration.name` - The configuration name to use when also setting
the `create_inactive` flag.

== Optional parameters:

* `create_inactive` - Create an AKV configuration without enabling it.
This flag is set to "false" by default.

== Related ONTAP commands

* `security key-manager external azure enable`
* `security key-manager external azure create-config`
* `security key-manager external azure update-config`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|create_inactive
|boolean
|query
|False
a|Indicates whether to create an active or inactive configuration.

* Introduced in: 9.16
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.
```

```
|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_certificate
|string
a|PKCS12 Certificate used by the application to prove its identity to AKV.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|client_secret
|string
a|Secret used by the application to prove its identity to AKV.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier of AKV key encryption key.
```

```
|name
|string
a|Name of the deployed AKV that will be used by ONTAP for storing keys.

* example: https://kmip-akv-keyvault.vault.azure.net/
* format: uri
* Introduced in: 9.8
* readCreate: 1
* x-nullable: true


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
```

```
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
a|AKV host subdomain.


|===


.Example request
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "authentication_method": "client_secret",
  "azure_reachability": {
    "code": "346758",
    "message": "AKV service is not reachable from all nodes - reason."
  },
  "client_certificate": "<CERTIFICATE-CONTENT>",
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
  "client_secret": "abcdef",
  "configuration": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "key_id":
"https://keyvault1.vault.azure.net/keys/key1/12345678901234567890123456789
012",
  "name": "https://kmip-akv-keyvault.vault.azure.net/",
  "oauth_host": "login.microsoftonline.com",
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "skip_verification": "",
  "state": {
    "code": "346758",
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
```

```
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vault_host": "vault.azure.net",
  "verify_host": "",
  "verify_ip": ""
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 3735553
| Failed to create self-signed certificate.

| 3735664
| The specified key size is not supported in FIPS mode.

| 3735665
| The specified hash function is not supported in FIPS mode.

| 3735700
| The specified key size is not supported.

| 52559972
| The certificates start date is later than the current date.

| 65537500
| A key manager has already been configured for this SVM.

| 65537504
| Internal error. Failed to store configuration in internal database.

| 65537505
| One or more volume encryption keys of the given SVM are stored on a key
manager configured for the admin SVM.

| 65537506
| AKV is not supported in MetroCluster configurations.

| 65537512
| AKV cannot be configured for the given SVM as not all nodes in the
cluster can enable the Azure Key Vault feature.

| 65537514
| Failed to check if the Azure Key Vault feature is enabled.

| 65537518
```

```
| Failed to find an interface with Cluster role.

| 65537523
| Invalid key ID format. Example key ID format":"
"https://mykeyvault.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e
2a74".

| 65537526
| Failed to enable Azure Key Vault feature.

| 65537553
| Invalid vault name format. Example vault name format":"
"https://mykeyvault.vault.azure.net".

| 65537567
| No authentication method provided.

| 65537573
| Invalid client certificate.

| 65537589
| The specified configuration.name already exists on the given SVM.

| 65537592
| The configuration.name field requires an ECV of 9.14.0 or greater.

| 65537593
| The create_inactive flag requires an effective cluster version of 9.14.0
or greater.

| 65537594
| The configuration.name field is required when the create_inactive flag
is set to true.

| 65537595
| The configuration.name field can only be specified when the
create_inactive flag is set to true.

| 65538902
| The configuration.name is reserved for use by the system.

| 65538903
| The configuration.name field cannot be an empty string.

| 65539704
| The key manager cannot be configured because the SVM has NAE volumes.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===



[#_links]
[.api-collapsible-fifth-title]
_links
[#azure_reachability]
[.api-collapsible-fifth-title]
azure_reachability

Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV service is
reachable from all nodes in the cluster.


|message
|string
a|Error message set when reachability is false.


|reachable
|boolean
a|Set to true when the AKV service is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
```

```
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the AKV wrapped internal key is available cluster
wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|boolean
a|Set to true when an AKV wrapped internal key is present on all nodes of
the cluster.
```

```
|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV wrapped key
is available on all nodes in the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#azure_key_vault]
[.api-collapsible-fifth-title]
azure_key_vault
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.


|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_certificate
|string
a|PKCS12 Certificate used by the application to prove its identity to AKV.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|client_secret
|string
a|Secret used by the application to prove its identity to AKV.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.
```

```
|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier of AKV key encryption key.


|name
|string
a|Name of the deployed AKV that will be used by ONTAP for storing keys.

* example: https://kmip-akv-keyvault.vault.azure.net/
* format: uri
* Introduced in: 9.8
* readCreate: 1
* x-nullable: true


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
```

```
a|Proxy port.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.
```

```
|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
a|AKV host subdomain.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID89cd4c04706a9385454c62f60f8cfe52]]
```

```
= Re-key the external key in the key hierarchy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/azure-key-vaults/{azure_key_vault.uuid}/rekey-external`#

*Introduced In:* 9.11

Rekeys the external key in the key hierarchy for an SVM with an AKV
configuration.

== Required properties

* `key_id` - Key identifier of the new AKV key encryption key.

== Related ONTAP commands

* `security key-manager external azure rekey-external`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|azure_key_vault.uuid
|string
|path
|True
a|UUID of the existing AKV configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
```

```
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier of the AKV key encryption key.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "key_id":
"https://keyvault1.vault.azure.net/keys/key1/12345678901234567890123456789
012"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537120
| Azure Key Vault is not configured for the given SVM.

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
```

```
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#azure_key_vault_key]
[.api-collapsible-fifth-title]
azure_key_vault_key

[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier of the AKV key encryption key.
```

```
|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
```

```
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDab7c54d8072335ef6c2e823a85a3a354]]
= Delete an AKV configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/azure-key-vaults/{uuid}`#

*Introduced In:* 9.8

Deletes an AKV configuration.

== Related ONTAP commands

* `security key-manager external azure disable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|AKV UUID


|return_timeout
|integer
|query
```

```
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536242
| One or more self-encrypting drives are assigned an authentication key.

| 65536243
| Cannot determine authentication key presence on one or more self-
encrypting drives.

| 65536817
| Internal error. Failed to determine if key manager is safe to disable.

| 65536827
| Internal error. Failed to determine if the given SVM has any encrypted
```

```
volumes.

| 65536834
| Internal error. Failed to get existing key-server details for the given
SVM.

| 65536867
| Volume encryption keys (VEK) for one or more encrypted volumes are
stored on the key manager configured for the given SVM.

| 65536883
| Internal error. Volume encryption key is missing for a volume.

| 65536884
| Internal error. Volume encryption key is invalid for a volume.

| 65536924
| Cannot remove key manager that still contains one or more NSE
authentication keys.

| 65537120
| Azure Key Vault is not configured for the given SVM.

| 196608080
| One or more nodes in the cluster have the root volume encrypted using
NVE (NetApp Volume Encryption).

| 196608301
| Internal error. Failed to get encryption type.

| 196608305
| NAE aggregates found in the cluster.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID2e8da22f64faa89d2695ffd649d02b0a]]
```

```
= Retrieve the AKV configuration for an SVM specified by the UUID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/azure-key-vaults/{uuid}`#

*Introduced In:* 9.8

Retrieves the AKV configuration for the SVM specified by the UUID.

== Related ONTAP commands

* `security key-manager external azure show`
* `security key-manager external azure check`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|AKV UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.


|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
```

```
a|Key Identifier of AKV key encryption key.


|name
|string
a|Name of the deployed AKV that will be used by ONTAP for storing keys.

* example: https://kmip-akv-keyvault.vault.azure.net/
* format: uri
* Introduced in: 9.8
* readCreate: 1
* x-nullable: true


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_port
|integer
a|Proxy port.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
```

```
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
a|AKV host subdomain.


|verify_host
|boolean
a|Verify the identity of the AKV host name.


|verify_ip
```

```
|boolean
a|Verify the identity of the AKV IP address.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_method": "client_secret",
  "azure_reachability": {
    "code": "346758",
    "message": "AKV service is not reachable from all nodes - reason."
  },
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
  "configuration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
```

```
  ],
  "key_id":
"https://keyvault1.vault.azure.net/keys/key1/12345678901234567890123456789
012",
  "name": "https://kmip-akv-keyvault.vault.azure.net/",
  "oauth_host": "login.microsoftonline.com",
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "skip_verification": "",
  "state": {
    "code": "346758",
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vault_host": "vault.azure.net",
  "verify_host": "",
  "verify_ip": ""
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description
```

```
| 65537549
| The Azure Key Vault Key Management Service is unreachable from one or
more nodes.

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
```

```
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#azure_reachability]
[.api-collapsible-fifth-title]
azure_reachability
```

```
Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV service is
reachable from all nodes in the cluster.


|message
|string
a|Error message set when reachability is false.


|reachable
|boolean
a|Set to true when the AKV service is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability
```

```
Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the AKV wrapped internal key is available cluster
```

```
wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|boolean
a|Set to true when an AKV wrapped internal key is present on all nodes of
the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV wrapped key
is available on all nodes in the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID082a6befc086236bbef2e5c0c20844eb]]
= Update the AKV configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/azure-key-vaults/{uuid}`#

*Introduced In:* 9.8

Updates the AKV configuration.

== Optional properties

* `client_secret` or `client_certificate` - New secret or new PKCS12
certificate used to prove the application's identity to the AKV.
```

```
* `proxy_type` - Type of proxy (http, https etc.) if proxy configuration
is used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `port` - Authorization server and vault port number.
* `oauth_host` - Open authorization server host name.
* `vault_host` - AKV host subdomain.
* `verify_host` - Verify the identity of the AKV host name.
* `verify_ip ` - Verify the identity of the AKV IP address.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `client_id` - Application (client) ID of the deployed Azure application
with appropriate access to an AKV.
* `tenant_id` - Directory (tenant) ID of the deployed Azure application
with appropriate access to an AKV.
* `skip_verification` - Skip the verification of the updated credentials,
set to true to bypass the verification. The default value is false.

== Related ONTAP commands

* `security key-manager external azure update-client-secret`
* `security key-manager external azure update-credentials`
* `security key-manager external azure update-config`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|AKV UUID


|return_timeout
|integer
|query
|False
```

```
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.


|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_certificate
|string
a|PKCS12 Certificate used by the application to prove its identity to AKV.
```

```
|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|client_secret
|string
a|Secret used by the application to prove its identity to AKV.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.
```

```
|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
a|AKV host subdomain.
```

```
|verify_host
|boolean
a|Verify the identity of the AKV host name.


|verify_ip
|boolean
a|Verify the identity of the AKV IP address.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "authentication_method": "client_secret",
  "azure_reachability": {
    "code": "346758",
    "message": "AKV service is not reachable from all nodes - reason."
  },
  "client_certificate": "<CERTIFICATE-CONTENT>",
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
  "client_secret": "abcdef",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "oauth_host": "login.microsoftonline.com",
  "port": 443,
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
```

```
  "skip_verification": "",
  "state": {
    "code": "346758",
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vault_host": "vault.azure.net",
  "verify_host": "",
  "verify_ip": ""
}
====

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537120
| Azure Key Vault is not configured for the given SVM.

| 65537504
| Internal error. Failed to store configuration in internal database.

| 65537517
| The field "client_secret" must be specified.

| 65537540
| Invalid client secret.
```

```
| 65537541
| No inputs were provided for the patch request.

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537573
| Invalid client certificate.

| 65537577
| The AKV certificate authentication method cannot be configured for the
given SVM as not all nodes in the cluster support the AKV certificate
authentication.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#azure_reachability]
[.api-collapsible-fifth-title]
azure_reachability

Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV service is
reachable from all nodes in the cluster.


|message
|string
a|Error message set when reachability is false.


|reachable
|boolean
a|Set to true when the AKV service is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration
```

```
Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability
```

```
Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the AKV wrapped internal key is available cluster
```

```
wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|boolean
a|Set to true when an AKV wrapped internal key is present on all nodes of
the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if AKV wrapped key
is available on all nodes in the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#azure_key_vault]
[.api-collapsible-fifth-title]
azure_key_vault

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Authentication method for the AKV instance.


|azure_reachability
|link:#azure_reachability[azure_reachability]
a|Indicates whether or not the AKV service is reachable from all the nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|client_certificate
|string
```

```
a|PKCS12 Certificate used by the application to prove its identity to AKV.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.


|client_secret
|string
a|Secret used by the application to prove its identity to AKV.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|oauth_host
|string
a|Open authorization server host name.


|port
|integer
a|Authorization server and vault port number.


|proxy_host
|string
a|Proxy host.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port.
```

```
|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|skip_verification
|boolean
a|Set to true to skip the verification of the updated user credentials
when updating credentials. The default value is false.


|state
|link:#state[state]
a|Indicates whether or not the AKV wrapped internal key is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|uuid
|string
a|A unique identifier for the Azure Key Vault (AKV).


|vault_host
|string
```

```
a|AKV host subdomain.


|verify_host
|boolean
a|Verify the identity of the AKV host name.


|verify_ip
|boolean
a|Verify the identity of the AKV IP address.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID8e1735aa113df85f278b902b68c71249]]
= Re-key the internal key in the key hierarchy for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/azure-key-vaults/{uuid}/rekey-internal`#

*Introduced In:* 9.10

Rekeys the internal key in the key hierarchy for an SVM with an AKV
configuration.

== Related ONTAP commands

* `security key-manager external azure rekey-internal`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the existing AKV configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537120
| Azure Key Vault is not configured for the given SVM.

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
```

```
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537559
| There are no existing internal keys for the SVM. A rekey operation is
allowed for an SVM with one or more encryption keys.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDdf0599e958fef20679be7a0db2fb0221]]
= Restore keys for an SVM from a configured AKV
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/azure-key-vaults/{uuid}/restore`#

*Introduced In:* 9.10

Restore the keys for an SVM from a configured AKV.

== Related ONTAP commands

* `security key-manager external azure restore`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the existing AKV configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0
```

```
|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537120
| Azure Key Vault is not configured for the given SVM.

| 65537515
| Failed to restore keys on some nodes in the cluster.

| 65537544
| Missing wrapped top-level internal key protection key (KEK) from
internal database.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage Barbican KMS configurations

:leveloffset: +1


[[ID7b36cc6d51c1ddaebb5689596b9192af]]
= Manage Barbican KMS configurations
```

```
== Overview

Barbican Key Management Services (KMS) is a key management service that
provides a secure store for encryption keys. This feature
allows ONTAP to securely protect its encryption keys using Barbican KMS.
Before you can use Barbican KMS with ONTAP, you must provide ONTAP with
the necessary details to allow ONTAP to communicate with the deployed
Barbican application.
These details include the key ID URL, Keystone authentication URL, and the
application credentials ID and secret.
The property `barbican_reachability` is considered an advanced property
and is populated only when explicitly requested.

== Examples

=== Creating an inactive Barbican configuration for an SVM

The example Barbican configuration is created for a specific SVM but is
not enabled.
Note the _return_records=true_ query parameter can be used to return the
newly created key-manager keystore configuration.

----

# The API:
POST /api/security/barbican-kms

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/barbican-kms' -H 'accept:
application/hal+json' -H 'Content-Type: application/json' -d '{ "svm": {
"name": "barbican_svm" }, "configuration": {"name": "myConfiguration"},
"application_cred_id": "app1", "application_cred_secret": "secret1",
"key_id": "https://sample.keyid.com:9311/v1/secrets/5c610a4f-ea97-44b5-
8682-f4daeafa9647/", "keystone_url":
"https://sample.keystone.com:5000/v3/auth/tokens" }'

# The response:
{
"job": {
  "uuid": "5a208fb0-fa58-11ef-8c9f-005056bbeee5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/5a208fb0-fa58-11ef-8c9f-005056bbeee5"
    }
```

```
  }
}
}
----

'''

=== Listing all Barbican configurations

The following example shows how to retrieve a list of all created Barbican
configurations.

----

# The API:
GET /api/security/barbican-kms

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/barbican-kms'

# The response:
{
"records": [
  {
    "configuration": {
      "name": "myConfiguration",
      "uuid": "5a134975-fa58-11ef-8c9f-005056bbeee5"
    },
    "uuid": "5a134975-fa58-11ef-8c9f-005056bbeee5",
    "_links": {
      "self": {
        "href": "/api/security/barbican-kms/5a134975-fa58-11ef-8c9f-
005056bbeee5"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/barbican-kms"
  }
}
}
----
```

```
'''

=== Retrieving a specific Barbican configuration

The following example shows how to retrieve information for a specific
Barbican configuration.

----

# The API:
GET /api/security/barbican-kms/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/barbican-kms/5a134975-fa58-
11ef-8c9f-005056bbeee5?fields=*'

# The response:
{
"configuration": {
  "name": "myConfiguration",
  "uuid": "5a134975-fa58-11ef-8c9f-005056bbeee5"
},
"uuid": "5a134975-fa58-11ef-8c9f-005056bbeee5",
"scope": "svm",
"svm": {
  "uuid": "ec8e0954-fa10-11ef-8c9f-005056bbeee5",
  "name": "barbican_svm"
},
"key_id": "https://sample.keyid.com:9311/v1/secrets/5c610a4f-ea97-44b5-
8682-f4daeafa9647/",
"application_cred_id": "app1",
"keystone_url": "https://sample.keystone.com:5000/v3/auth/tokens",
"proxy_type": "https",
"proxy_host": "",
"proxy_port": 0,
"proxy_username": "",
"verify": true,
"verify_host": true,
"enabled": false,
"timeout": 10,
"_links": {
  "self": {
    "href": "/api/security/barbican-kms/5a134975-fa58-11ef-8c9f-
005056bbeee5"
  }
}
```

```
}
----

'''

=== Retrieving an advanced property for a specific Barbican configuration

The following example shows how to retrieve an advanced property for a
specific Barbican configuration.

----

# The API:
GET /api/security/barbican-kms/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/barbican-kms/f72098a2-e908-
11ea-bd56-005056bb4222?fields=barbican_reachability'

# The response:
{
"configuration": {
  "name": "myConfiguration",
  "uuid": "f72098a2-e908-11ea-bd56-005056bb4222"
},
"uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
"barbican_reachability": {
  "reachable": true,
  "message": "",
  "code": "0"
},
"_links": {
  "self": {
    "href": "/api/security/barbican-kms/f72098a2-e908-11ea-bd56-
005056bb4222"
  }
}
}
----

'''

=== Updating the application credentials ID and secret for a specific
Barbican configuration

The following example shows how to update the application credentials for
a specific Barbican configuration.
```

```
----

# The API:
PATCH /api/security/barbican-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/barbican-kms/f72098a2-e908-
11ea-bd56-005056bb4222/' -d '{"application_cred_id": "app345",
"application_cred_secret" : "secret"}'

# The response:
{
"job": {
  "uuid": "921ee66a-fa5a-11ef-8c9f-005056bbeee5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/921ee66a-fa5a-11ef-8c9f-005056bbeee5"
    }
  }
}
}
----

'''

=== Enabling a Barbican configuration

The newly created Barbican configuration is inactive by default. Use the
REST API PATCH method "/api/security/key-stores/\{uuid}" to enable the
configuration.

'''

=== Restoring keys

The following example shows how to restore keys for a specific Barbican
configuration.

----

# The API:
POST /api/security/barbican-kms/{uuid}/restore

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/barbican-kms/f72098a2-e908-
11ea-bd56-005056bb4222/restore'
```

```
# The response:
{
"job": {
  "uuid": "e89e7283-2c49-11f0-831c-005056bb33a5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e89e7283-2c49-11f0-831c-005056bb33a5"
    }
  }
}
}
----

'''

=== Rekey the internal key

The following example shows how to rekey the internal key based on a
specific Barbican configuration.

----

# The API:
POST /api/security/barbican-kms/{uuid}/rekey-internal

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/barbican-kms/f72098a2-e908-
11ea-bd56-005056bb4222/rekey-internal'

# The response:
{
"job": {
  "uuid": "921ee66a-fa5a-11ef-8c9f-005056bbeee5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/921ee66a-fa5a-11ef-8c9f-005056bbeee5"
    }
  }
}
}
----

'''
```

```
[[ID1f75d5f9f6f8c5814ff4e0eeec1e2c32]]
= Retrieve Barbican KMS configurations for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/barbican-kms`#

*Introduced In:* 9.17

Retrieves Barbican KMS configurations for all SVMs.

== Related ONTAP commands

* `security key-manager external barbican show`
* `security key-manager external barbican check`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|keystone_url
|string
|query
|False
a|Filter by keystone_url


|uuid
|string
|query
|False
a|Filter by uuid


|timeout
|integer
|query
|False
```

```
a|Filter by timeout


|proxy_type
|string
|query
|False
a|Filter by proxy_type


|enabled
|boolean
|query
|False
a|Filter by enabled


|proxy_port
|integer
|query
|False
a|Filter by proxy_port


|barbican_reachability.message
|string
|query
|False
a|Filter by barbican_reachability.message


|barbican_reachability.code
|string
|query
|False
a|Filter by barbican_reachability.code


|barbican_reachability.reachable
|boolean
|query
|False
a|Filter by barbican_reachability.reachable


|configuration.uuid
|string
```

```
|query
|False
a|Filter by configuration.uuid


|configuration.name
|string
|query
|False
a|Filter by configuration.name


|state.cluster_state
|boolean
|query
|False
a|Filter by state.cluster_state


|state.message
|string
|query
|False
a|Filter by state.message


|state.code
|string
|query
|False
a|Filter by state.code


|key_id
|string
|query
|False
a|Filter by key_id


|proxy_host
|string
|query
|False
a|Filter by proxy_host
```

```
|scope
|string
|query
|False
a|Filter by scope


|verify
|boolean
|query
|False
a|Filter by verify


|proxy_username
|string
|query
|False
a|Filter by proxy_username


|application_cred_id
|string
|query
|False
a|Filter by application_cred_id


|verify_host
|boolean
|query
|False
a|Filter by verify_host


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#barbican[barbican]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "application_cred_id": "63e3cb77f84f42b7a0395a3efb7636f9",
      "barbican_reachability": {
        "code": "346758",
        "message": "Barbican KMS is not reachable from all nodes -
<reason>.",
        "reachable": ""
      },
      "configuration": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      },
      "key_id": "https://172.29.58.184:9311/v1/secrets/5c610a4f-ea97-44b5-
8682-f4daeafa9647",
      "keystone_url": "https://keystoneip:5000/v3/auth/tokens",
      "proxy_host": "proxy.eng.com",
      "proxy_port": 1234,
      "proxy_type": "http",
      "proxy_username": "proxyuser",
      "scope": "string",
      "state": {
        "cluster_state": "",
        "code": "346758",
        "message": "Top-level internal key encryption key is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "timeout": 60,
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
```

```
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65539801
| The Barbican Key Management Service is unreachable from one or more
nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
```

```
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#barbican_reachability]
[.api-collapsible-fifth-title]
barbican_reachability

Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Code corresponding to the error message. Returns 0 if Barbican KMS is
reachable from all nodes in the cluster.


|message
|string
a|Set to the appropriate error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Barbican KMS is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===
```

```
[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when an SVM-KEK is present on all nodes of the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if the SVM-KEK is
available on all nodes in the cluster.


|message
|string
a|Error message returned when there's no SVM-KEK availability on the
cluster.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#barbican]
[.api-collapsible-fifth-title]
barbican

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
```

```
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier URL of the Barbican KMS key encryption key. Must be an
HTTPS URL.


|keystone_url
|string
a|Keystone URL for the access token. Must be an HTTPS URL.


|proxy_host
|string
a|Proxy host name.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
```

```
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|Connection timeout in seconds.


|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4c0b23e3e8fb4f181ee8c2f5ac3aafe6]]
= Create a Barbican KMS configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/barbican-kms`#

*Introduced In:* 9.17

Creates a Barbican KMS configuration for the specified SVM.

== Required properties

* `configuration.name` - Name for the new Barbican configuration.
* `svm.uuid` or `svm.name` - Existing SVM in which to create a Barbican
KMS.
* `key_id` - Barbican key URL.
* `keystone_url` - Keystone authentication URL.
* `application_cred_id` - Keystone authentication application ID with
access to the Barbican KMS.
* `application_cred_secret`- Application credentials secret to
authenticate the application credentials ID with Keystone.

== Optional properties

* `proxy_type` - Type of proxy (http/https) if proxy configuration is
used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `verify` - Verify the identity of the Barbican KMS?
* `verify_host` - Verify the identity of the Barbican KMS host name?
```

```
* `timeout` - Connection timeout in seconds.

== Related ONTAP commands

* `security key-manager external barbican create-config`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|application_cred_secret
|string
a|Keystone application credentials secret required to access the specified
Barbican KMS. It is not audited.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
```

```
a|Key Identifier URL of the Barbican KMS key encryption key. Must be an
HTTPS URL.


|keystone_url
|string
a|Keystone URL for the access token. Must be an HTTPS URL.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
```

```
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|Connection timeout in seconds.


|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "application_cred_id": "63e3cb77f84f42b7a0395a3efb7636f9",
  "application_cred_secret": "secret",
  "barbican_reachability": {
    "code": "346758",
    "message": "Barbican KMS is not reachable from all nodes - <reason>.",
    "reachable": ""
  },
  "configuration": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
```

```
  },
  "key_id": "https://172.29.58.184:9311/v1/secrets/5c610a4f-ea97-44b5-
8682-f4daeafa9647",
  "keystone_url": "https://keystoneip:5000/v3/auth/tokens",
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "state": {
    "cluster_state": "",
    "code": "346758",
    "message": "Top-level internal key encryption key is unavailable on
the following nodes with the associated reasons: Node: node1. Reason: No
volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timeout": 60,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
}
====

== Response
```


Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1115127
| Package MT_EK_MGMT is not licensed in the cluster.

| 65539805
| Failed to create a Barbican Key Management Service configuration because
the "application-cred-secret" field is invalid or was not provided.

| 65539806
| Barbican Key Management Service (KMS) cannot be configured because not
all nodes in the cluster are running an effective cluster version of 9.17
or later to support the Barbican KMS feature.

| 65539807
| Barbican Key Management Service is not supported for the admin SVM.

| 65539808
| Failed to create the specified Barbican Key Management Service
configuration because a Barbican Key Management Service configuration with
the same name already exists.

| 65539809
| Barbican Key Management Service is not supported in MetroCluster
configurations.

| 65539810
| Failed to create a Barbican Key Management Service configuration because
the "key-id" field is invalid or was not provided. The URI scheme must be
HTTPS.

| 65539811
| Failed to create a Barbican Key Management Service configuration because
the "keystone-url" field is invalid or was not provided. The URI scheme
must be HTTPS.

| 65539835
| Failed to create a Barbican Key Management Service configuration because
```

```
the "application-cred-id" field cannot be empty.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#barbican_reachability]
[.api-collapsible-fifth-title]
barbican_reachability

Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns 0 if Barbican KMS is
reachable from all nodes in the cluster.


|message
|string
a|Set to the appropriate error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Barbican KMS is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
```

```
a|Keystore UUID.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when an SVM-KEK is present on all nodes of the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if the SVM-KEK is
available on all nodes in the cluster.


|message
|string
a|Error message returned when there's no SVM-KEK availability on the
cluster.


|===


[#svm]
[.api-collapsible-fifth-title]
```

```
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#barbican]
[.api-collapsible-fifth-title]
barbican

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|application_cred_secret
|string
```

```
a|Keystone application credentials secret required to access the specified
Barbican KMS. It is not audited.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier URL of the Barbican KMS key encryption key. Must be an
HTTPS URL.


|keystone_url
|string
a|Keystone URL for the access token. Must be an HTTPS URL.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
```

```
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|Connection timeout in seconds.


|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
```

```
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0fa2340c57c8358bd05a6f62539d9255]]
= Retrieve the Barbican KMS configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/barbican-kms/{uuid}`#

*Introduced In:* 9.17

Retrieves the Barbican KMS configuration for the SVM specified by the
UUID.

== Related ONTAP commands

* `security key-manager external barbican show`
* `security key-manager external barbican check`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Barbican KMS UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|key_id
|string
a|Key Identifier URL of the Barbican KMS key encryption key. Must be an
HTTPS URL.


|keystone_url
|string
a|Keystone URL for the access token. Must be an HTTPS URL.
```

```
|proxy_host
|string
a|Proxy host name.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|timeout
|integer
a|Connection timeout in seconds.
```

```
|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application_cred_id": "63e3cb77f84f42b7a0395a3efb7636f9",
  "barbican_reachability": {
    "code": "346758",
    "message": "Barbican KMS is not reachable from all nodes - <reason>.",
    "reachable": ""
  },
  "configuration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "key_id": "https://172.29.58.184:9311/v1/secrets/5c610a4f-ea97-44b5-
8682-f4daeafa9647",
  "keystone_url": "https://keystoneip:5000/v3/auth/tokens",
```

```
  "proxy_host": "proxy.eng.com",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "state": {
    "cluster_state": "",
    "code": "346758",
    "message": "Top-level internal key encryption key is unavailable on
the following nodes with the associated reasons: Node: node1. Reason: No
volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "timeout": 60,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65539801
| The Barbican Key Management Service is unreachable from one or more
nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#barbican_reachability]
[.api-collapsible-fifth-title]
barbican_reachability

Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Code corresponding to the error message. Returns 0 if Barbican KMS is
reachable from all nodes in the cluster.


|message
|string
a|Set to the appropriate error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Barbican KMS is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.
```

```
|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when an SVM-KEK is present on all nodes of the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if the SVM-KEK is
available on all nodes in the cluster.


|message
|string
a|Error message returned when there's no SVM-KEK availability on the
cluster.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDbb9d1dc45bab260feaaebfc81a16db60]]
= Update the Barbican KMS configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/barbican-kms/{uuid}`#

*Introduced In:* 9.17
```

```
Updates the Barbican KMS configuration.

== Optional properties

* `application_cred_id` - New credentials used to verify the application's
identity to the Barbican KMS. You must provide both `application_cred_id`
and `application_cred_secret` to update the credentials.
* `application_cred_secret` - New credentials secret used to verify the
application's identity to the Barbican KMS. You must provide both
`application_cred_id` and `application_cred_secret` to update the
credentials.
* `proxy_type` - Type of proxy (http/https) if proxy configuration is
used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `verify` - Verify the identity of the Barbican KMS?
* `verify_host` - Verify the identity of the Barbican KMS host name?
* `timeout` - Connection timeout in seconds.

== Related ONTAP commands

* `security key-manager external barbican update-credentials`
* `security key-manager external barbican update-config`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Barbican KMS UUID


|return_timeout
```

```
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|application_cred_secret
|string
a|Keystone application credentials secret required to access the specified
Barbican KMS. It is not audited.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|timeout
|integer
a|Connection timeout in seconds.


|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "application_cred_id": "63e3cb77f84f42b7a0395a3efb7636f9",
  "application_cred_secret": "secret",
  "barbican_reachability": {
    "code": "346758",
    "message": "Barbican KMS is not reachable from all nodes - <reason>.",
    "reachable": ""
  },
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
```

```
  "scope": "string",
  "state": {
    "cluster_state": "",
    "code": "346758",
    "message": "Top-level internal key encryption key is unavailable on
the following nodes with the associated reasons: Node: node1. Reason: No
volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "timeout": 60,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65539223
| Failed to check the health of the Barbican Key Management Service.

| 65539814
| Failed to update the Barbican Key Management Service configuration
because the "application-cred-secret" field is invalid or was not
provided.

| 65539815
| Failed to update the Barbican Key Management Service configuration
because the "application-cred-id" field is invalid or was not provided.

| 65539832
| Specified Barbican Key Management Service configuration does not exist.

| 65539841
| Both "application_cred_id" and "application_cred_secret" are required.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#barbican_reachability]
[.api-collapsible-fifth-title]
barbican_reachability

Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns 0 if Barbican KMS is
reachable from all nodes in the cluster.


|message
|string
a|Set to the appropriate error message when 'reachable' is false.


|reachable
|boolean
```

```
a|Set to true if the Barbican KMS is reachable from all nodes of the
cluster.


|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#state]
[.api-collapsible-fifth-title]
state

Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when an SVM-KEK is present on all nodes of the cluster.


|code
|string
a|Code corresponding to the status message. Returns a 0 if the SVM-KEK is
available on all nodes in the cluster.


|message
|string
a|Error message returned when there's no SVM-KEK availability on the
cluster.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#barbican]
[.api-collapsible-fifth-title]
barbican

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_cred_id
|string
a|Keystone application credentials ID required to access the specified
Barbican KMS.


|application_cred_secret
|string
a|Keystone application credentials secret required to access the specified
Barbican KMS. It is not audited.


|barbican_reachability
|link:#barbican_reachability[barbican_reachability]
a|Indicates whether the Barbican KMS is reachable from all nodes in the
cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET request or an instance GET request unless it is explicitly
requested using the field's query parameter or GET for all advanced
properties is enabled.


|enabled
|boolean
a|Indicates whether the configuration is enabled.
```

```
|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Indicates whether or not the SVM key encryption key (KEK) is available
cluster wide.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|timeout
|integer
a|Connection timeout in seconds.
```

```
|uuid
|string
a|A unique identifier of the Barbican KMS.


|verify
|boolean
a|Verify the identity of the Barbican KMS.


|verify_host
|boolean
a|Verify the identity of the Barbican KMS host name.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[ID5f308a9f146bfc4910762dab9c447036]]
= Rekey the internal key in the key hierarchy for an SVM with a Barbican
KMS configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/barbican-kms/{uuid}/rekey-internal`#

*Introduced In:* 9.17

Rekeys the internal key in the key hierarchy for an SVM with a Barbican
KMS configuration.

== Related ONTAP commands

* `security key-manager external barbican rekey-internal`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the existing Barbican KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
```

```
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536205
| Internal error. Failed to generate SVM key in the kernel.

| 65536882
| Internal error. UUID is missing.

| 65536883
```

```
| Internal error. Volume encryption key is missing.

| 65536884
| Internal error. Volume encryption key is invalid.

| 65536889
| Internal error. Volume has invalid encryption blob.

| 65536973
| Internal error. Volume DSID is missing.

| 65537533
| Internal error. Failed to unwrap SVM key encryption key.

| 65537547
| One or more volume encryption keys (VEKs) for this data SVM's encrypted
volumes are stored in the key manager configured for the admin SVM. Use
the REST API POST method "/api/security/key-
managers/{source.svm.uuid}/migrate" to migrate these keys from the admin
SVM's key manager to this data SVM's key manager before running the rekey
operation.

| 65537556
| ONTAP is unable to encrypt or decrypt because the configured external
key manager is in a blocked state. Possible reasons for a blocked state
include the top-level external key protection key is not found, disabled
or has insufficient privileges. Resolve the external key manager key
issues at the key manager's portal before running the rekey operation.

| 65537565
| Internal error. Failed to update the VDEK blob.

| 65537610
| Rekey cannot be performed while the enabled keystore configuration is
being switched. If a previous attempt to switch the keystore configuration
failed, or was interrupted, the system will continue to prevent rekeying.
Use the REST API PATCH method "/api/security/key-stores/\{uuid}" to re-run
and complete the operation.

| 65539200
| The key custodian not provided with an external key value store.

| 65539201
| Failed to encrypt.

| 65539202
| Failed to decrypt.
```

```
| 65539416
| Internal error. Failed to parse the key value store response.

| 65539436
| Rekey cannot be performed while the enabled keystore configuration is
being initialized. Wait until the keystore is in the active state, and
rerun the rekey operation.

| 65539437
| Rekey cannot be performed while the enabled keystore configuration is
being disabled.

| 65539817
| The Barbican Key Management Service is not configured.

| 65539834
| New Key ID not found.

| 65539836
| No Key Management Service configured.

| 65539844
| Failed to import the SVM-KEK.

| 196608088
| Internal error. Failed to get encryption operation status.

| 196608352
| Internal error. The encryption blob was not found in the RDB table.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID97ebbda5083e41be93b698581b82a2e0]]
= Restore the keys for an SVM from a configured Barbican KMS

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/barbican-kms/{uuid}/restore`#

*Introduced In:* 9.17

Restores the keys for an SVM from a configured Barbican KMS.

== Related ONTAP commands

* `security key-manager external barbican restore`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the existing Barbican KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.
```

```
* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65539817
| The Barbican Key Management Service is not configured.

| 65539842
| Failed to restore the SVM-KEK on at least one node.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


[[ID6726091c959f4ed420689fae694e6784]]
= Create a certificate signing request

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/certificate-signing-request`#

*Introduced In:* 9.8
```

```
This API generates a Certificate Signing Request(CSR) and a private key
pair. A CSR is a message sent securely to a certificate authority (CA) via
any electronic media to apply for a digital identity certificate. This is
a general utility API for users to generate a CSR.

== Required properties

* `subject_name` - Subject details of the certificate, including but not
limited to the common name. Can be omitted if subject_alternatives is
specified.
* `subject_alternatives` - Subject Alternate Name (SAN) extensions. Only
required if a common name is not specified using subject_name.

== Recommended optional properties

* `security_strength` - Key size of the certificate, in bits. Specifying a
stronger security strength in bits is recommended when creating a
certificate. A value of at least 128 is necessary when using EC algorithm
if the certificate is to be used in the context of TLSv1.3.
* `hash_function` -  Hashing function.
* `algorithm` - Asymmetric algorithm. Algorithm used to generate a
public/private key pair when creating a certificate.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `security_strength` - _112_
* `hash_function` - _sha256_
* `algorithm` - _rsa_

== Related ONTAP commands

* `security certificate generate-csr`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|Asymmetric Encryption Algorithm.


|csr
|string
a|A Certificate Signing Request (CSR) provided to a CA for obtaining a CA-
signed certificate.


|extended_key_usages
|array[string]
a|A list of extended key usage extensions.


|generated_private_key
|string
a|Private key generated for the CSR.


|hash_function
|string
a|Hashing function.
```

```
|key_usages
|array[string]
a|A list of key usage extensions.


|security_strength
|integer
a|Security strength of the certificate in bits.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_name
|string
a|Subject name details of the certificate. The format is a list of comma
separated key=value pairs.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "algorithm": "string",
  "csr": "string",
  "extended_key_usages": [
    "string"
  ],
  "generated_private_key": "string",
  "hash_function": "string",
  "key_usages": [
    "string"
  ],
  "security_strength": 112,
  "subject_alternatives": {
    "dns": [
      "*.example.com"
    ],
    "email": [
      "abc@example.com"
    ],
```

```
    "ip": [
      "10.225.34.10"
    ],
    "uri": [
      "http://example.com"
    ]
  },
  "subject_name": "C=US,O=NTAP,CN=test.domain.com"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|Asymmetric Encryption Algorithm.


|csr
|string
a|A Certificate Signing Request (CSR) provided to a CA for obtaining a CA-
signed certificate.


|extended_key_usages
|array[string]
a|A list of extended key usage extensions.


|generated_private_key
|string
a|Private key generated for the CSR.


|hash_function
|string
a|Hashing function.
```

```
|key_usages
|array[string]
a|A list of key usage extensions.


|security_strength
|integer
a|Security strength of the certificate in bits.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_name
|string
a|Subject name details of the certificate. The format is a list of comma
separated key=value pairs.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "algorithm": "string",
  "csr": "string",
  "extended_key_usages": [
    "string"
  ],
  "generated_private_key": "string",
  "hash_function": "string",
  "key_usages": [
    "string"
  ],
  "security_strength": 112,
  "subject_alternatives": {
    "dns": [
      "*.example.com"
    ],
    "email": [
      "abc@example.com"
```

```
    ],
    "ip": [
      "10.225.34.10"
    ],
    "uri": [
      "http://example.com"
    ]
  },
  "subject_name": "C=US,O=NTAP,CN=test.domain.com"
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 3735554
| Certificate signing request failed.

| 3735664
| Key size is not supported in FIPS mode.

| 3735665
| Hash function is not supported in FIPS mode.

| 3735700
| Key size is not supported.

| 3735713
| Security strength bits length is not supported.

| 3735714
| Security strength bits length is not supported in FIPS mode.

| 3735715
| Certificate creation requires a common name or SAN extensions.

| 3735741
| Key size is not applicable with the EC encryption algorithm.

| 3735750
```

```
| Cannot supply "critical" as the only value in a list of extension
values.

| 52560173
| Hash function is not supported for digital signatures.

| 52560423
| Failed to read the relative distinguished names.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#subject_alternatives]
[.api-collapsible-fifth-title]
subject_alternatives

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|dns
|array[string]
a|A list of DNS names for Subject Alternate name extension.


|email
|array[string]
a|A list of email addresses for Subject Alternate name extension


|ip
|array[string]
a|A list of IP addresses for Subject Alternate name extension.


|uri
|array[string]
a|A list of URIs for Subject Alternate name extension.


|===


[#certificate_signing_request]
[.api-collapsible-fifth-title]
certificate_signing_request

[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|Asymmetric Encryption Algorithm.


|csr
|string
a|A Certificate Signing Request (CSR) provided to a CA for obtaining a CA-
signed certificate.
```

```
|extended_key_usages
|array[string]
a|A list of extended key usage extensions.


|generated_private_key
|string
a|Private key generated for the CSR.


|hash_function
|string
a|Hashing function.


|key_usages
|array[string]
a|A list of key usage extensions.


|security_strength
|integer
a|Security strength of the certificate in bits.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_name
|string
a|Subject name details of the certificate. The format is a list of comma
separated key=value pairs.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



= Manage security certificates

:leveloffset: +1


[[ID560f455ca413d094e90ac9d0bfb7bf9c]]
= Manage security certificates



:doctype: book

== Overview

This API displays security certificate information and manages the
certificates in ONTAP.

== Installing certificates in ONTAP

The security certificates GET request retrieves all of the certificates in
the cluster.

== Examples

=== Retrieving all certificates installed in the cluster with their
common-names

----

# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?fields=common_name" -H "accept:
application/hal+json"

# The response:
{
"records": [
```

```
  {
    "svm": {
      "name": "vs0"
    },
    "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
    "common_name": "vs0",
    "_links": {
      "self": {
        "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-
005056b482fc"
      }
    }
  },
  {
    "uuid": "1941e048-8ac1-11e8-9058-005056b482fc",
    "common_name": "ROOT",
    "_links": {
      "self": {
        "href": "/api/security/certificates/1941e048-8ac1-11e8-9058-
005056b482fc"
      }
    }
  },
  {
    "uuid": "5a3a77a8-892d-11e8-b7da-005056b482fc",
    "common_name": "cert_name",
    "_links": {
      "self": {
        "href": "/api/security/certificates/5a3a77a8-892d-11e8-b7da-
005056b482fc"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/security/certificates?fields=common_name"
  }
}
}
----

'''

=== Retrieving all certificates installed at cluster-scope with their
```

```
common-names

'''

----

# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?scope=cluster&fields=common_name" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "1941e048-8ac1-11e8-9058-005056b482fc",
    "scope": "cluster",
    "common_name": "ROOT",
    "_links": {
      "self": {
        "href": "/api/security/certificates/1941e048-8ac1-11e8-9058-
005056b482fc"
      }
    }
  },
  {
    "uuid": "5a3a77a8-892d-11e8-b7da-005056b482fc",
    "scope": "cluster",
    "common_name": "cert_name",
    "_links": {
      "self": {
        "href": "/api/security/certificates/5a3a77a8-892d-11e8-b7da-
005056b482fc"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/certificates?scope=cluster&fields=common_name"
  }
}
```

```
}
----

'''

=== Retrieving all certificates installed on a specific SVM with their
common-names

'''

----

# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?svm.name=vs0&fields=common_name" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "svm": {
      "name": "vs0"
    },
    "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
    "common_name": "vs0",
    "_links": {
      "self": {
        "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-
005056b482fc"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/certificates?svm.name=vs0&fields=common_name"
  }
}
}
----
```

```
'''

=== Retrieving a certificate using its UUID for all fields

'''

----

# The API:
/api/security/certificates/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/security/certificates/dad2363b-8ac0-
11e8-9058-005056b482fc?fields=*" -H "accept: application/hal+json"

# The response:
{
"svm": {
  "uuid": "d817293c-8ac0-11e8-9058-005056b482fc",
  "name": "vs0"
},
"uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
"scope": "svm",
"type": "server",
"common_name": "vs0",
"serial_number": "15428D45CF81CF56",
"ca": "vs0",
"hash_function": "sha256",
"key_size": 2048,
"expiry_time": "2019-07-18T15:29:14-04:00",
"public_certificate": "<CERTIFICATE-CONTENT>",
"_links": {
  "self": {
    "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-
005056b482fc"
  }
}
}
----

=== Creating a certificate in a cluster

These certificates can be used to help administrators enable certificate-
based authentication and to enable SSL-based communication to the cluster.

----
```

```
# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"common_name\": \"TEST-SERVER\",  \"type\": \"server\"  }"
----

=== Installing a certificate in a cluster

These certificates can be used to help administrators enable certificate-
based authentication and to enable-SSL based communication to the cluster.

----

# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{ \"type\":
\"server_ca\", \"public_certificate\": \"<CERTIFICATE-CONTENT>\"  }"
----

'''

=== Installing a certificate on a specific SVM

'''

----

# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{   \"svm\" : {
\"name\" : \"vs0\" }, \"type\": \"server_ca\", \"public_certificate\":
\"<CERTIFICATE-CONTENT>\"  }"
----

'''

=== Installing a CA-signed certificate on a specific SVM
```

```
'''

----

# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{   \"svm\" : {
\"name\" : \"vs0\" }, \"type\": \"server\", \"public_certificate\":
\"<CERTIFICATE-CONTENT>\", \"intermediate_certificates\": [\"<CERTIFICATE-
CONTENT>\",\"<CERTIFICATE-CONTENT>\"]}"
----

'''

=== Deleting a certificate using its UUID

'''

----

# The API:
/api/security/certificates/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/certificates/dad2363b-8ac0-
11e8-9058-005056b482fc?fields=*" -H "accept: application/hal+json"
----

=== Signing a new certificate signing request using an existing CA
certificate UUID

Once you have created a certificate of type "root_ca", you can use that
certificate to act as a local Certificate Authority to sign new
certificate signing requests. The following example signs a new
certificate signing request using an existing CA certificate UUID. If
successful, the API returns a signed certificate.

----

# The API:
/api/security/certificates/{ca.uuid}/sign

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates/253add53-8ac9-
```

```
11e8-9058-005056b482fc/sign" -H  "accept: application/json" -H  "Content-
Type: application/json" -d "{ \"signing_request\": \"<CERTIFICATE-
CONTENT>\",  \"hash_function\": \"sha256\"}"

# The response:
{
"public_certificate": "<CERTIFICATE-CONTENT>"
}
----

=== Generate a new Certificate Signing Request (CSR)

----

# The API:
/api/security/certificate-signing-request

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificate-signing-request"
-H  "accept: application/json" -H  "Content-Type: application/json" -d "{
\"algorithm\": \"rsa\", \"extended_key_usage\": [\"serverauth\"],
\"hash_function\": \"sha256\",  \"key_usage\": [\"digitalsignature\"],
\"security_strength\": \"112\", \"subject_alternatives\": { \"dns\": [
\"*.example.com\", \"*.example1.com\" ], \"email\": [\"abc@example.com\",
\"abc@example1.com\"], \"ip\": [\"10.225.34.223\", \"10.225.34.224\"],
\"uri\": [\"http://example.com\", \"http://example1.com\"] },
\"subject_name\": \"C=US,O=NTAP,CN=test.domain.com\"}"
{
"csr": "-----BEGIN CERTIFICATE REQUEST-----\n-----END CERTIFICATE
REQUEST-----\n",
"generated_private_key": "-----BEGIN PRIVATE KEY-----\n-----END PRIVATE
KEY-----\n"
}
----

'''

----

### Download and install a certificate from the Azure Key Vault.
----

= The API:

/api/security/certificates

= The call:
```

```
curl -X POST "https://+++<mgmt-ip>+++/api/security/certificates" -H
"accept: application/json" -H "Content-Type: application/json" -d
"{\"svm\": {\"name\" : \"vs0\"}, \"name\" : \"vs0-client-cert\", \"type\"
: \"client\", \"azure\": { \"key_vault\" :
\"https://example.vault.azure.net\", \"client_id\" : \"12345678-abcd-1234-
12ad-dfasdffgfdaaa\", \"tenant_id\" : \"12345678-abcd-abcd-test-
720ef604b100\", \"client_secret\" : \"clientSecretString\",
\"verify_host\": false }}" { "job": { "uuid": "be8d45cb-1d41-11ee-9725-
005056ae0f31", "_links": { "self": { "href": "/api/cluster/jobs/be8d45cb-
1d41-11ee-9725-005056ae0f31" } } } }+++</mgmt-ip>+++

----


'''
----

[discrete]
=== Creating a root-ca certificate in a cluster when DCN cluster is
configured.

----

# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H  "accept:
application/json" -H  "Content-Type: application/json" -d "{
\"common_name\": \"TEST-ROOT-CA\",  \"type\": \"root_ca\"  }"

# The response:
{
"job": {
  "uuid": "be8d45cb-1d41-11ee-9725-005056ae0f31",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/be8d45cb-1d41-11ee-9725-005056ae0f31"
    }
  }
}
}
----

'''
```

```
[discrete]
=== Deleting a CA certificate using its UUID when DCN cluster is
configured.

'''

----

# The API:
/api/security/certificates/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/certificates/dad2363b-8ac0-
11e8-9058-005056b482fc?fields=*" -H "accept: application/hal+json"

# The response:
{
"job": {
  "uuid": "be8d45cb-1d41-11ee-9725-005056ae0f31",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/be8d45cb-1d41-11ee-9725-005056ae0f31"
    }
  }
}
}
----

'''





[[ID d2d7ea434d0d736bec260798fc845925]]
= Retrieve security certificates

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/certificates`#

*Introduced In:* 9.6

Retrieves security certificates.

== Related ONTAP commands

* `security certificate show`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|subject_key_identifier
|string
|query
|False
a|Filter by subject_key_identifier

* Introduced in: 9.8


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|ca
|string
|query
|False
a|Filter by ca

* maxLength: 256
* minLength: 1


|serial_number
```

```
|string
|query
|False
a|Filter by serial_number

* maxLength: 40
* minLength: 1


|common_name
|string
|query
|False
a|Filter by common_name


|scope
|string
|query
|False
a|Filter by scope


|hash_function
|string
|query
|False
a|Filter by hash_function


|expiry_time
|string
|query
|False
a|Filter by expiry_time


|public_certificate
|string
|query
|False
a|Filter by public_certificate

* Introduced in: 9.8


|name
```

```
|string
|query
|False
a|Filter by name

* Introduced in: 9.8


|type
|string
|query
|False
a|Filter by type


|subject_alternatives.uri
|string
|query
|False
a|Filter by subject_alternatives.uri

* Introduced in: 9.15


|subject_alternatives.ip
|string
|query
|False
a|Filter by subject_alternatives.ip

* Introduced in: 9.15


|subject_alternatives.email
|string
|query
|False
a|Filter by subject_alternatives.email

* Introduced in: 9.15


|subject_alternatives.dns
|string
|query
|False
a|Filter by subject_alternatives.dns
```

```
* Introduced in: 9.15


|authority_key_identifier
|string
|query
|False
a|Filter by authority_key_identifier

* Introduced in: 9.8


|key_size
|integer
|query
|False
a|Filter by key_size


|uuid
|string
|query
|False
a|Filter by uuid

* Introduced in: 9.8


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
```

```
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records
```

```
|records
|array[link:#security_certificate[security_certificate]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
      "azure": {
        "verify_host": true
      },
      "ca": "string",
      "common_name": "test.domain.com",
      "expiry_time": "2030-01-25 11:20:13 +0000",
      "hash_function": "string",
      "key_size": 512,
      "name": "string",
      "public_certificate": "<CERTIFICATE-CONTENT>",
      "scope": "string",
      "serial_number": "string",
      "subject_alternatives": {
        "dns": [
```

```
          "*.example.com"
        ],
        "email": [
          "abc@example.com"
        ],
        "ip": [
          "10.225.34.10"
        ],
        "uri": [
          "http://example.com"
        ]
      },
      "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#proxy]
[.api-collapsible-fifth-title]
proxy
[#azure]
[.api-collapsible-fifth-title]
azure

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|proxy
|link:#proxy[proxy]
a|

|===


[#subject_alternatives]
[.api-collapsible-fifth-title]
subject_alternatives

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dns
|array[string]
a|A list of DNS names for Subject Alternate name extension.


|email
|array[string]
a|A list of email addresses for Subject Alternate name extension


|ip
|array[string]
a|A list of IP addresses for Subject Alternate name extension.


|uri
|array[string]
a|A list of URIs for Subject Alternate name extension.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_certificate]
[.api-collapsible-fifth-title]
security_certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authority_key_identifier
|string
a|Provides the key identifier of the issuing CA certificate that signed
the SSL certificate.


|azure
```

```
|link:#azure[azure]
a|

|ca
|string
a|Certificate authority


|common_name
|string
a|FQDN or custom common name. Provide on POST when creating a self-signed
certificate.


|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
time format. Can be provided on POST if creating self-signed certificate.
The expiration time range is between 1 day to 10 years.


|hash_function
|string
a|Hashing function. Can be provided on POST when creating a self-signed
certificate. Hash functions md5 and sha1 are not allowed on POST.


|key_size
|integer
a|Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048,
3072. Can be provided on POST if creating self-signed certificate with a
minimum permissible value of 2048.


|name
|string
a|Certificate name or name of the certificate to be downloaded from the
Azure Key Vault (AKV). If not provided in POST, a unique name specific to
the SVM is automatically generated.


|public_certificate
|string
a|Public key Certificate in PEM format. If this is not provided in POST, a
self-signed certificate is created.
```

```
|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|serial_number
|string
a|Serial number of certificate.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_key_identifier
|string
a|Provides the key identifier used to identify the public key in the SSL
certificate.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of Certificate. The following types are supported:

* client - a certificate and its private key used by an SSL client in
ONTAP.
* server - a certificate and its private key used by an SSL server in
ONTAP.
* client_ca - a Certificate Authority certificate used by an SSL server in
ONTAP to verify an SSL client certificate.
* server_ca - a Certificate Authority certificate used by an SSL client in
ONTAP to verify an SSL server certificate.
* root_ca - a self-signed certificate used by ONTAP to sign other
certificates by acting as a Certificate Authority.
* enum: ["client", "server", "client_ca", "server_ca", "root_ca"]
* Introduced in: 9.6
* x-nullable: true


|uuid
|string
```

```
a|Unique ID that identifies a certificate.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID441f7cbcbfaf6d4e13f9fec6a3b62cff]]
= Create or install security certificates

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/certificates`#

*Introduced In:* 9.6

Creates or installs a certificate or downloads a certificate from Azure
Key Vault (AKV) and installs it on the ONTAP cluster.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create or install
the certificate.
* `common_name` - Common name of the certificate. Required when creating a
certificate.
* `type` - Type of certificate.
* `public_certificate` - Public key certificate in PEM format. Required
when installing a certificate.
* `private_key` - Private key certificate in PEM format. Required when
installing a CA-signed certificate.

== Recommended optional properties

* `expiry_time` - Certificate expiration time. Specifying an expiration
time is recommended when creating a certificate.
```

```
* `key_size` - Key size of the certificate in bits. Specifying a strong
key size is recommended when creating a certificate.
* `name` - Unique certificate name per SVM or the name of the certificate
in AKV, required for downloading AKV certificates. If one is not provided,
it is automatically generated.

== AKV required properties for downloading a certificate

* `azure.key_vault` - URI of the Azure Key Vault.
* `azure.client_id` - Application (client) ID of the deployed Azure
application with appropriate access to an AKV.
* `azure.tenant_id` - Directory (tenant) ID of the deployed Azure
application with appropriate access to an AKV.
* `azure.client_secret` - Secret used by the application to prove its
identity to AKV.
* `azure.client_certificate` - PKCS12 certificate used by the application
to prove its identity to AKV.

== AKV optional properties for downloading a certificate

* `azure.oauth_host` - Open authorization server host name.
* `azure.proxy.type` - Type of proxy (http, https etc.) if proxy
configuration is used.
* `azure.proxy.host` - Proxy hostname if proxy configuration is used.
* `azure.proxy.port` - Proxy port number if proxy configuration is used.
* `azure.proxy.username` - Proxy username if proxy configuration is used.
* `azure.proxy.password` - Proxy password if proxy configuration is used.
* `azure.timeout` - AKV connection timeout in seconds.
* `azure.verify_host` - Verify the identity of the AKV host name.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `key_size` - _2048_
* `expiry_time` - _P365DT_
* `hash_function` - _sha256_

== Related ONTAP commands

* `security certificate create`
* `security certificate install`
* `security certificate azure-install`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|authority_key_identifier
|string
a|Provides the key identifier of the issuing CA certificate that signed
the SSL certificate.


|azure
|link:#azure[azure]
a|

|ca
|string
a|Certificate authority


|common_name
|string
```

```
a|FQDN or custom common name. Provide on POST when creating a self-signed
certificate.


|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
time format. Can be provided on POST if creating self-signed certificate.
The expiration time range is between 1 day to 10 years.


|hash_function
|string
a|Hashing function. Can be provided on POST when creating a self-signed
certificate. Hash functions md5 and sha1 are not allowed on POST.


|intermediate_certificates
|array[string]
a|Chain of intermediate Certificates in PEM format. Only valid in POST
when installing a certificate.


|key_size
|integer
a|Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048,
3072. Can be provided on POST if creating self-signed certificate with a
minimum permissible value of 2048.


|name
|string
a|Certificate name or name of the certificate to be downloaded from the
Azure Key Vault (AKV). If not provided in POST, a unique name specific to
the SVM is automatically generated.


|private_key
|string
a|Private key Certificate in PEM format. Only valid for create when
installing a CA-signed certificate. This is not audited.


|public_certificate
|string
a|Public key Certificate in PEM format. If this is not provided in POST, a
self-signed certificate is created.
```

```
|serial_number
|string
a|Serial number of certificate.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_key_identifier
|string
a|Provides the key identifier used to identify the public key in the SSL
certificate.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of Certificate. The following types are supported:

* client - a certificate and its private key used by an SSL client in
ONTAP.
* server - a certificate and its private key used by an SSL server in
ONTAP.
* client_ca - a Certificate Authority certificate used by an SSL server in
ONTAP to verify an SSL client certificate.
* server_ca - a Certificate Authority certificate used by an SSL client in
ONTAP to verify an SSL server certificate.
* root_ca - a self-signed certificate used by ONTAP to sign other
certificates by acting as a Certificate Authority.
* enum: ["client", "server", "client_ca", "server_ca", "root_ca"]
* Introduced in: 9.6
* x-nullable: true


|uuid
|string
a|Unique ID that identifies a certificate.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
  "azure": {
    "client_certificate": "PEM Cert",
    "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
    "client_secret": "abcdef",
    "key_vault": "https://kmip-akv-keyvault.vault.azure.net/",
    "oauth_host": "login.microsoftonline.com",
    "proxy": {
      "host": "proxy.eng.com",
      "password": "proxypassword",
      "port": 1234,
      "type": "string",
      "username": "proxyuser"
    },
    "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
    "timeout": 25
  },
  "ca": "string",
  "common_name": "test.domain.com",
  "expiry_time": "2030-01-25 11:20:13 +0000",
  "hash_function": "string",
  "intermediate_certificates": [
    "<CERTIFICATE-CONTENT>"
  ],
  "key_size": 512,
  "name": "string",
  "private_key": "-----BEGIN PRIVATE KEY-----\\nprivate-key\\n-----END
PRIVATE KEY-----\\n",
  "public_certificate": "<CERTIFICATE-CONTENT>",
  "serial_number": "string",
  "subject_alternatives": {
    "dns": [
      "*.example.com"
    ],
    "email": [
      "abc@example.com"
    ],
    "ip": [
```

```
      "10.225.34.10"
    ],
    "uri": [
      "http://example.com"
    ]
  },
  "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#security_certificate[security_certificate]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
```

```
    {
      "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
      "azure": {
        "client_certificate": "PEM Cert",
        "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
        "client_secret": "abcdef",
        "key_vault": "https://kmip-akv-keyvault.vault.azure.net/",
        "oauth_host": "login.microsoftonline.com",
        "proxy": {
          "host": "proxy.eng.com",
          "password": "proxypassword",
          "port": 1234,
          "type": "string",
          "username": "proxyuser"
        },
        "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
        "timeout": 25
      },
      "ca": "string",
      "common_name": "test.domain.com",
      "expiry_time": "2030-01-25 11:20:13 +0000",
      "hash_function": "string",
      "intermediate_certificates": [
        "<CERTIFICATE-CONTENT>"
      ],
      "key_size": 512,
      "name": "string",
      "private_key": "-----BEGIN PRIVATE KEY-----\\nprivate-key\\n-----END
PRIVATE KEY-----\\n",
      "public_certificate": "<CERTIFICATE-CONTENT>",
      "serial_number": "string",
      "subject_alternatives": {
        "dns": [
          "*.example.com"
        ],
        "email": [
          "abc@example.com"
        ],
        "ip": [
          "10.225.34.10"
        ],
        "uri": [
          "http://example.com"
        ]
      },
```

```
      "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```


Status: Default


```
ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|3735645  +
```

```
//end row
//start row
|Cannot specify a value for serial. It is generated automatically.
//end row
//start row
|
//end row
//start row
|3735622  +
//end row
//start row
|The certificate type is not supported.
//end row
//start row
|
//end row
//start row
|3735664  +
//end row
//start row
|The specified key size is not supported in FIPS mode.
//end row
//start row
|
//end row
//start row
|3735665  +
//end row
//start row
|The specified hash function is not supported in FIPS mode.
//end row
//start row
|
//end row
//start row
|3735553  +
//end row
//start row
|Failed to create self-signed Certificate.
//end row
//start row
|
//end row
//start row
|3735646  +
//end row
```

```
//start row
|Failed to store the certificates.
//end row
//start row
|
//end row
//start row
|3735693  +
//end row
//start row
|The certificate installation failed as private key was empty.
//end row
//start row
|
//end row
//start row
|3735618  +
//end row
//start row
|Cannot accept private key for server_ca or client_ca.
//end row
//start row
|
//end row
//start row
|52363365 +
//end row
//start row
|Failed to allocate memory.
//end row
//start row
|
//end row
//start row
|52559975 +
//end row
//start row
|Failed to read the certificate due to incorrect formatting.
//end row
//start row
|
//end row
//start row
|52363366 +
//end row
//start row
```

```
|Unsupported key type.
//end row
//start row
|
//end row
//start row
|52560123 +
//end row
//start row
|Failed to read the key due to incorrect formatting.
//end row
//start row
|
//end row
//start row
|52559972 +
//end row
//start row
|The certificates start date is later than the current date.
//end row
//start row
|
//end row
//start row
|52559976 +
//end row
//start row
|The certificate and private key do not match.
//end row
//start row
|
//end row
//start row
|52559973 +
//end row
//start row
|The certificate has expired.
//end row
//start row
|
//end row
//start row
|52363366 +
//end row
//start row
|Logic error: use of a dead object.
```

```
//end row
//start row
|
//end row
//start row
|3735696  +
//end row
//start row
|Intermediate certificates are not supported with client_ca and server_ca
type certificates.
//end row
//start row
|
//end row
//start row
|52559974 +
//end row
//start row
|The certificate is not supported in FIPS mode.
//end row
//start row
|
//end row
//start row
|3735676  +
//end row
//start row
|Cannot continue the installation without a value for the common name.
Since the subject field in the certificate is empty, the field
"common_name" must have a value to continue with the installation.
//end row
//start row
|
//end row
//start row
|3735558  +
//end row
//start row
|Failed to extract information about Common Name from the certificate.
//end row
//start row
|
//end row
//start row
|3735588  +
//end row
```

```
//start row
|The common name (CN) extracted from the certificate is not valid.
//end row
//start row
|
//end row
//start row
|3735632  +
//end row
//start row
|Failed to extract Certificate Authority Information from the certificate.
//end row
//start row
|
//end row
//start row
|3735700  +
//end row
//start row
|The specified key size is not supported.
//end row
//start row
|
//end row
//start row
|52560173 +
//end row
//start row
|The hash function is not supported for digital signatures.
//end row
//start row
|
//end row
//start row
|3735751  +
//end row
//start row
|Failed to authenticate and fetch the access token from Azure OAuth host.
//end row
//start row
|
//end row
//start row
|3735752  +
//end row
//start row
```

```
|Failed to extract the private key from the Azure Key Vault certificate.
//end row
//start row
|3735753  +
//end row
//start row
|Unsupported content_type in the Azure secrets response.
//end row
//start row
|3735754  +
//end row
//start row
|Internal error. Failed to parse the JSON response from Azure Key Vault.
//end row
//start row
|3735755  +
//end row
//start row
|REST call to Azure failed.
//end row
//start row
|3735756  +
//end row
//start row
|Invalid client certificate.
//end row
//start row
|3735757  +
//end row
//start row
|Internal error. Failed to generate client assertion.
//end row
//start row
|3735762  +
//end row
//start row
|Provided Azure Key Vault configuration is incorrect.
//end row
//start row
|3735763  +
//end row
//start row
|Provided Azure Key Vault configuration is incomplete.
//end row
//start row
|3735764  +
```

```
//end row
//start row
|Request to Azure failed. Reason - Azure error code and Azure error
message.
//end row
|===
//end table


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#proxy]
[.api-collapsible-fifth-title]
proxy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|host
```

```
|string
a|Proxy host.


|password
|string
a|Proxy password. Password is not audited.


|port
|integer
a|Proxy port.


|type
|string
a|Proxy type.


|username
|string
a|Proxy username.


|===


[#azure]
[.api-collapsible-fifth-title]
azure

[cols=3*,options=header]
|===
|Name
|Type
|Description

|client_certificate
|string
a|PKCS12 certificate used by the application to prove its identity to AKV.


|client_id
|string
a|Application client ID of the deployed Azure application with appropriate
access to an AKV.
```

```
|client_secret
|string
a|Secret used by the application to prove its identity to AKV.


|key_vault
|string
a|URI of the deployed AKV that is used by ONTAP for storing keys.

* example: https://kmip-akv-keyvault.vault.azure.net/
* format: uri
* x-ntap-createOnly: true
* Introduced in: 9.14
* x-nullable: true


|oauth_host
|string
a|Open authorization server host name.


|proxy
|link:#proxy[proxy]
a|

|tenant_id
|string
a|Directory (tenant) ID of the deployed Azure application with appropriate
access to an AKV.


|timeout
|integer
a|AKV connection timeout, in seconds. The allowed range is between 0 to 30
seconds.


|verify_host
|boolean
a|Verify the identity of the AKV host name. By default, verify_host is set
to true.


|===
```

```
[#subject_alternatives]
[.api-collapsible-fifth-title]
subject_alternatives

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dns
|array[string]
a|A list of DNS names for Subject Alternate name extension.


|email
|array[string]
a|A list of email addresses for Subject Alternate name extension


|ip
|array[string]
a|A list of IP addresses for Subject Alternate name extension.


|uri
|array[string]
a|A list of URIs for Subject Alternate name extension.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_certificate]
[.api-collapsible-fifth-title]
security_certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authority_key_identifier
|string
a|Provides the key identifier of the issuing CA certificate that signed
the SSL certificate.


|azure
|link:#azure[azure]
a|

|ca
|string
a|Certificate authority


|common_name
|string
a|FQDN or custom common name. Provide on POST when creating a self-signed
certificate.


|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
```

```
time format. Can be provided on POST if creating self-signed certificate.
The expiration time range is between 1 day to 10 years.


|hash_function
|string
a|Hashing function. Can be provided on POST when creating a self-signed
certificate. Hash functions md5 and sha1 are not allowed on POST.


|intermediate_certificates
|array[string]
a|Chain of intermediate Certificates in PEM format. Only valid in POST
when installing a certificate.


|key_size
|integer
a|Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048,
3072. Can be provided on POST if creating self-signed certificate with a
minimum permissible value of 2048.


|name
|string
a|Certificate name or name of the certificate to be downloaded from the
Azure Key Vault (AKV). If not provided in POST, a unique name specific to
the SVM is automatically generated.


|private_key
|string
a|Private key Certificate in PEM format. Only valid for create when
installing a CA-signed certificate. This is not audited.


|public_certificate
|string
a|Public key Certificate in PEM format. If this is not provided in POST, a
self-signed certificate is created.


|serial_number
|string
a|Serial number of certificate.
```

```
|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_key_identifier
|string
a|Provides the key identifier used to identify the public key in the SSL
certificate.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of Certificate. The following types are supported:

* client - a certificate and its private key used by an SSL client in
ONTAP.
* server - a certificate and its private key used by an SSL server in
ONTAP.
* client_ca - a Certificate Authority certificate used by an SSL server in
ONTAP to verify an SSL client certificate.
* server_ca - a Certificate Authority certificate used by an SSL client in
ONTAP to verify an SSL server certificate.
* root_ca - a self-signed certificate used by ONTAP to sign other
certificates by acting as a Certificate Authority.
* enum: ["client", "server", "client_ca", "server_ca", "root_ca"]
* Introduced in: 9.6
* x-nullable: true


|uuid
|string
a|Unique ID that identifies a certificate.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID38964e58ebecbd854f06342fe6964c46]]
= Sign security certificates

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/certificates/{ca.uuid}/sign`#

*Introduced In:* 9.6

Signs a certificate.

== Required properties

* `signing_request` - Certificate signing request to be signed by the
given certificate authority.

== Recommended optional properties

* `expiry_time` - Certificate expiration time. Specifying an expiration
time for a signed certificate is recommended.
* `hash_function` - Hashing function. Specifying a strong hashing function
is recommended when signing a certificate.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `expiry_time` - _P365DT_
* `hash_function` - _sha256_

== Related ONTAP commands

* `security certificate sign`
This API is used to sign a certificate request using a pre-existing self-
signed root certificate. The self-signed root certificate acts as a
certificate authority within its scope and maintains the records of its
signed certificates.
```

```
The root certificate can be created for a given SVM or for the cluster
using [`POST security/certificates`].


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|ca.uuid
|string
|path
|True
a|UUID of the existing certificate authority certificate


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
time format. The allowed expiration time range is between 1 day to 10
years.
```

```
|hash_function
|string
a|Hashing function


|signing_request
|string
a|Certificate signing request to be signed by the given certificate
authority. Request should be in X509 PEM format.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "expiry_time": "P1DT2H3M4S or '2030-01-25T11:20:13Z'",
  "hash_function": "string",
  "signing_request": "<CERTIFICATE-CONTENT>"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|public_certificate
|string
a|CA signed public key Certificate


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "public_certificate": "string"
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 3735628
| Failed to use CA certificate for signing.

| 3735665
| The specified hash function is not supported in FIPS mode.

| 52559974
| The certificate is not supported in FIPS mode.

| 3735626
| Failed to generate signed Certificate.

| 3735558
```

```
| Failed to extract information about Common Name from the certificate.

| 3735588
| The common name (CN) extracted from the certificate is not valid.

| 3735632
| Failed to extract Certificate Authority Information from the
certificate.

| 3735629
| Failed to sign the certificate because Common Name of signing
certificate and Common Name of CA certificate are same.

| 3735630
| Failed to sign the certificate because expiry date of signing
certificate exceeds the expiry date of CA certificate.

| 3735701
| Invalid expiration period. The allowed range for expiration time is
between 1 and 3652 days.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#security_certificate_sign]
[.api-collapsible-fifth-title]
security_certificate_sign

[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
time format. The allowed expiration time range is between 1 day to 10
years.
```

```
|hash_function
|string
a|Hashing function


|signing_request
|string
a|Certificate signing request to be signed by the given certificate
authority. Request should be in X509 PEM format.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf5599ae7a10ba4a941acbf65df2908dd]]
= Delete security certificates

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/certificates/{uuid}`#

*Introduced In:* 9.6

Deletes a security certificate.

== Related ONTAP commands

* `security certificate delete`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Certificate UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 3735644
| Cannot delete server-chain certificate. Reason: There is a corresponding
server certificate for it.

| 3735679
| Cannot delete pre-installed server_ca certificates through REST. Use CLI
or ZAPI.

| 3735650
| Deleting this client_ca certificate directly is not supported. Delete
the corresponding root-ca certificate using type `root_ca` to delete the
root, client, and server certificates.

| 3735627
| Deleting this server_ca certificate directly is not supported. Delete
the corresponding root-ca certificate using type `root_ca` to delete the
root, client, and server certificates.
```

```
| 3735589
| Cannot delete certificate.

| 3735590
| Cannot delete certificate. Failed to remove SSL configuration for the
certificate.

| 3735683
| Cannot remove this certificate while external key manager is configured.

| 3735681
| Cannot delete preinstalled `server-ca` certificates. Use the CLI to
complete the operation.

| 52560272
| The certificate could not be removed due to being in use by one or more
subsystems.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
```

```
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf67656e52361ab32d6994350062a234c]]
= Retrieve security certificates

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/certificates/{uuid}`#

*Introduced In:* 9.6

Retrieves security certificates.

== Related ONTAP commands

* `security certificate show`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Certificate UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authority_key_identifier
|string
a|Provides the key identifier of the issuing CA certificate that signed
the SSL certificate.


|azure
|link:#azure[azure]
a|
```

```
|ca
|string
a|Certificate authority


|common_name
|string
a|FQDN or custom common name. Provide on POST when creating a self-signed
certificate.


|expiry_time
|string
a|Certificate expiration time, in ISO 8601 duration format or date and
time format. Can be provided on POST if creating self-signed certificate.
The expiration time range is between 1 day to 10 years.


|hash_function
|string
a|Hashing function. Can be provided on POST when creating a self-signed
certificate. Hash functions md5 and sha1 are not allowed on POST.


|key_size
|integer
a|Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048,
3072. Can be provided on POST if creating self-signed certificate with a
minimum permissible value of 2048.


|name
|string
a|Certificate name or name of the certificate to be downloaded from the
Azure Key Vault (AKV). If not provided in POST, a unique name specific to
the SVM is automatically generated.


|public_certificate
|string
a|Public key Certificate in PEM format. If this is not provided in POST, a
self-signed certificate is created.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
```

```
"cluster".


|serial_number
|string
a|Serial number of certificate.


|subject_alternatives
|link:#subject_alternatives[subject_alternatives]
a|

|subject_key_identifier
|string
a|Provides the key identifier used to identify the public key in the SSL
certificate.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of Certificate. The following types are supported:

* client - a certificate and its private key used by an SSL client in
ONTAP.
* server - a certificate and its private key used by an SSL server in
ONTAP.
* client_ca - a Certificate Authority certificate used by an SSL server in
ONTAP to verify an SSL client certificate.
* server_ca - a Certificate Authority certificate used by an SSL client in
ONTAP to verify an SSL server certificate.
* root_ca - a self-signed certificate used by ONTAP to sign other
certificates by acting as a Certificate Authority.
* enum: ["client", "server", "client_ca", "server_ca", "root_ca"]
* Introduced in: 9.6
* x-nullable: true


|uuid
|string
a|Unique ID that identifies a certificate.
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
  "azure": {
    "verify_host": true
  },
  "ca": "string",
  "common_name": "test.domain.com",
  "expiry_time": "2030-01-25 11:20:13 +0000",
  "hash_function": "string",
  "key_size": 512,
  "name": "string",
  "public_certificate": "<CERTIFICATE-CONTENT>",
  "scope": "string",
  "serial_number": "string",
  "subject_alternatives": {
    "dns": [
      "*.example.com"
    ],
    "email": [
      "abc@example.com"
    ],
    "ip": [
      "10.225.34.10"
    ],
    "uri": [
      "http://example.com"
    ]
  },
  "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
  "svm": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#proxy]
[.api-collapsible-fifth-title]
proxy
[#azure]
```

```
[.api-collapsible-fifth-title]
azure

[cols=3*,options=header]
|===
|Name
|Type
|Description

|proxy
|link:#proxy[proxy]
a|

|===


[#subject_alternatives]
[.api-collapsible-fifth-title]
subject_alternatives

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dns
|array[string]
a|A list of DNS names for Subject Alternate name extension.


|email
|array[string]
a|A list of email addresses for Subject Alternate name extension


|ip
|array[string]
a|A list of IP addresses for Subject Alternate name extension.


|uri
|array[string]
a|A list of URIs for Subject Alternate name extension.


|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage cluster network security configuration

:leveloffset: +1


[[ID86d712cf8a4256e782c50c1739dee3cf]]
= Manage cluster network security configuration



== Overview

You can use the security cluster-network API endpoints to modify the
cluster network security configuration.
The following operations are supported:

* GET to retrieve the cluster network security status: GET
security/cluster-network
* PATCH to update the cluster network security configuration: PATCH
security/cluster-network





[[IDd16b5e38ae8fee2a4fe561bb1c0edb7f]]
= Retrieve the cluster network security configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/cluster-network`#

*Introduced In:* 9.18

Retrieves the cluster network security configuration.

== Related ONTAP commands

* 'security cluster-network show'


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|enabled
|boolean
a|Indicates whether cluster network security is enabled.


|mode
|string
a|The cluster network security mode.


|status
|string
a|The status of the cluster network security configuration.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "mode": "string",
  "status": "ENABLING | DISABLING | READY"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4bb271fb27aafc4d7ba42ebfe52d8529]]
= Update the cluster network security configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

```
block]#`/security/cluster-network`#

*Introduced In:* 9.18

Updates the cluster network security configuration.

== Optional properties

* 'enabled' - Indicates whether cluster network security is enabled.
* 'mode' - The cluster network security mode.

== Related ONTAP commands

* 'security cluster-network modify'


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether cluster network security is enabled.


|mode
|string
a|The cluster network security mode.


|status
|string
a|The status of the cluster network security configuration.


|===


.Example request
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "mode": "string",
  "status": "ENABLING | DISABLING | READY"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 70320128
| Modifying the cluster network security configuration requires an
effective cluster version of ONTAP 9.18.1 or later.

| 70320131
| Cluster network security does not support MetroCluster configurations.

| 70320132
| The cluster has an out of quorum node, wait until all nodes are in
quorum before modifying the cluster network security configuration.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
```

```
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#security_cluster_network]
[.api-collapsible-fifth-title]
security_cluster_network

Manages the cluster network security configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether cluster network security is enabled.


|mode
|string
a|The cluster network security mode.


|status
|string
a|The status of the cluster network security configuration.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage cluster network security certificates

:leveloffset: +1


[[ID471b94ed7518d17c490272013d105729]]
= Manage cluster network security certificates



== Overview

You can use the security cluster-network certificate API endpoints to
view and modify the certificate configuration for cluster network
security.
The following operations are supported:

* GET to retrieve the certificate configuration for cluster network
security: GET security/cluster-network/certificate
* PATCH to update the certificate configuration for cluster network
security for a given node: PATCH security/cluster-network/certificate
* POST to specify the certificate configuration for cluster network
security for a given node: POST security/cluster-network/certificate
* DELETE to remove the certificate configuration for cluster network
security for a given node: DELETE security/cluster-network/certificate
```

```
[[ID0d896b2b2483fd13b4d72594768da923]]
= Retrieve cluster network security certificate configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/cluster-network/certificates`#

*Introduced In:* 9.18

Retrieves the certificate configuration used for cluster network security.

== Related ONTAP commands

* 'security cluster-network certificate show'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|certificate.name
|string
|query
|False
```

```
a|Filter by certificate.name


|certificate.uuid
|string
|query
|False
a|Filter by certificate.uuid


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#security_certificate_reference]
[.api-collapsible-fifth-title]
security_certificate_reference

Security certificate object reference
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#records]
[.api-collapsible-fifth-title]
records

Manages the cluster network security certificate configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd643b14f09c665ec1a28b92e43c44293]]
= Create a cluster network security certificate configuration
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/cluster-network/certificates`#

*Introduced In:* 9.18

Specifies the certificate configuration for cluster network security for a
given node.

== Required properties

* 'node' - The node to which the certificate will be assigned.
* 'name' - The certificate name.

== Related ONTAP commands

* 'security cluster-network certificate create'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference
```

```
|node
|link:#node_reference[node_reference]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 70320128
| Modifying the cluster network security configuration requires an
effective cluster version of ONTAP 9.18.1 or later.

| 70320132
| The cluster has an out of quorum node, wait until all nodes are in
quorum before modifying the cluster network security configuration.

| 70320137
| The specified certificate was not found on the admin SVM
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links
[#security_certificate_reference]
[.api-collapsible-fifth-title]
security_certificate_reference

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#security_cluster_network_certificates]
[.api-collapsible-fifth-title]
security_cluster_network_certificates

Manages the cluster network security certificate configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4acf9ca1b91677ab1b8c651e20a51dd1]]
= Delete a cluster network security certificate configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/cluster-network/certificates/{node.uuid}`#

*Introduced In:* 9.18
```

```
Deletes the certificate configuration for cluster network security for a
given node.

== Required properties

* `node: Node UUID`

== Related ONTAP commands

* 'security cluster-network certificate delete'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 70320128
| Modifying the cluster network security configuration requires an
effective cluster version of ONTAP 9.18.1 or later.

| 70320132
| The cluster has an out of quorum node, wait until all nodes are in
quorum before modifying the cluster network security configuration.

| 70320137
| The specified certificate was not found on the admin SVM

| 70320150
| The certificate table is currently in use. Removing a certificate entry
is not allowed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#security_certificate_reference]
[.api-collapsible-fifth-title]
security_certificate_reference

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf8e1fd5ca4aa531f699f4b26ed732085]]
= Retrieve cluster network security certificate configuration for a node

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/cluster-network/certificates/{node.uuid}`#

*Introduced In:* 9.18

Retrieves the certificate configuration for cluster network security for a
given node.

== Related ONTAP commands

* 'security cluster-network certificate show'


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#security_certificate_reference]
[.api-collapsible-fifth-title]
security_certificate_reference

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID
```

```
|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID15fe1a72a86be6c9d1d54f3aa18ec9b4]]
= Update a cluster network security certificate configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/cluster-network/certificates/{node.uuid}`#
```

```
*Introduced In:* 9.18

Updates the certificate configuration for cluster network security for a
given node.

== Required properties

* `node: Node UUID`

== Related ONTAP commands

* 'security cluster-network certificate modify'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference
```

```
|node
|link:#node_reference[node_reference]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 70320128
| Modifying the cluster network security configuration requires an
effective cluster version of ONTAP 9.18.1 or later.

| 70320132
| The cluster has an out of quorum node, wait until all nodes are in
quorum before modifying the cluster network security configuration.
```

```
| 70320137
| The specified certificate was not found on the admin SVM

| 70320138
| The certificate table is currently in use. Removing a certificate entry
is not allowed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===



[#_links]
[.api-collapsible-fifth-title]
_links
[#security_certificate_reference]
[.api-collapsible-fifth-title]
security_certificate_reference

Security certificate object reference
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#security_cluster_network_certificates]
[.api-collapsible-fifth-title]
security_cluster_network_certificates

Manages the cluster network security certificate configuration.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#security_certificate_reference[security_certificate_reference]
a|Security certificate object reference


|node
|link:#node_reference[node_reference]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage external-role-mapping entries

:leveloffset: +1


[[ID78a496402fec778b67749ac1811df4c5]]
= Manage external-role-mapping entries



=== Overview
```

```
This API is used to map external identity provider roles to ONTAP roles.
The POST request creates an external role to ONTAP role mapping. Various
responses are shown in the examples below.

=== Examples

=== Creating an external role to ONTAP role mapping entry

The following output shows how to create an external role to ONTAP role
mapping entry.

----

# The API
/api/security/external-role-mappings

# The call:
curl -X POST "https://<mgmt-ip>/api/security/external-role-mappings" -d '{
"external_role": "Administrator", "ontap_role": { "name": "admin" },
"provider": "adfs", "comment": "Admin role" }'
----

=== Retrieving all external role to ONTAP role mapping entries

The following output shows all external role to ONTAP role mapping
entries.

----

# The API
/api/security/external-role-mappings

# The call:
curl -X GET "https://<mgmt-ip>/api/security/external-role-mappings"

# The response:
{
 "records": [
 {
   "external_role": "Administrator",
   "provider": "adfs"
 }
 ],
 "num_records": 1
}
----
```

```
[[ID938633ec8ea1529856cde7557d306f55]]
= Retrieve all external-role-mapping entries

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/external-role-mappings`#

*Introduced In:* 9.16

Retrieves all external-role-mapping entries.

== Related ONTAP commands

* `security login external-role-mapping show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|comment
|string
|query
|False
a|Filter by comment


|provider
|string
|query
|False
a|Filter by provider


|ontap_role.name
|string
```

```
|query
|False
a|Filter by ontap_role.name


|timestamp
|string
|query
|False
a|Filter by timestamp


|external_role
|string
|query
|False
a|Filter by external_role


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
```

```
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_external_role_mapping[security_external_role_mapping
]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "external_role": "string",
      "ontap_role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "provider": "string",
      "timestamp": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role

ONTAP role to map with the external identity role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Role name


|===


[#security_external_role_mapping]
[.api-collapsible-fifth-title]
security_external_role_mapping

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#self_link[self_link]
a|

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|external_role
|string
a|External Identity provider role.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|provider
|string
a|Type of the external identity provider.


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[ID1c02916894b6ba15e5c064fdd6b3df11]]
= Create an external-role-mapping entry

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/external-role-mappings`#

*Introduced In:* 9.16

Creates an external-role-mapping entry.

== Required properties

* `external_role`
* `provider`
* `ontap_role`

== Optional properties

* `comment`

== Related ONTAP commands

* `security login external-role-mapping create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:
```

```
|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|external_role
|string
a|External Identity provider role.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|provider
|string
a|Type of the external identity provider.


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


.Example request
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "comment": "string",
  "external_role": "string",
  "ontap_role": {
    "name": "admin"
  },
  "provider": "string",
  "timestamp": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#security_external_role_mapping[security_external_role_mapping
]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "comment": "string",
      "external_role": "string",
      "ontap_role": {
        "name": "admin"
```

```
      },
      "provider": "string",
      "timestamp": "string"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5636243
| Provided ONTAP role is not configured.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role

ONTAP role to map with the external identity role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Role name


|===
```

```
[#security_external_role_mapping]
[.api-collapsible-fifth-title]
security_external_role_mapping

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|external_role
|string
a|External Identity provider role.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|provider
|string
a|Type of the external identity provider.


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb8108b0117e9865e400d9d20d024a9e6]]
= Delete an external-role-mapping entry

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/external-role-mappings/{external_role}/{provider}`#

*Introduced In:* 9.16

Deletes an external-role-mapping entry.

== Related ONTAP commands

* `security login external-role-mapping delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|external_role
|string
|path
|True
a|

|provider
|string
```

```
|path
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0c06042c9e94e5c069fb903aed3766fe]]
= Retrieve an external-role-mapping entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/external-role-mappings/{external_role}/{provider}`#

*Introduced In:* 9.16

Retrieves an external-role-mapping entry.

== Related ONTAP commands

* `security login external-role-mapping show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|external_role
|string
|path
|True
a|External Identity provider role.


|provider
|string
|path
|True
a|Type of the external identity provider.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|external_role
|string
a|External Identity provider role.
```

```
|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|provider
|string
a|Type of the external identity provider.


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "external_role": "string",
  "ontap_role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "provider": "string",
  "timestamp": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#ontap_role]
```

```
[.api-collapsible-fifth-title]
ontap_role

ONTAP role to map with the external identity role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Role name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDef046116cc55e5d2126df61dda7214f1]]
= Update an external-role-mapping entry

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/external-role-mappings/{external_role}/{provider}`#

*Introduced In:* 9.16

Updates an external-role-mapping entry.
```

```
== Related ONTAP commands

* `security login external-role-mapping modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|external_role
|string
|path
|True
a|External Identity provider role.


|provider
|string
|path
|True
a|Type of the external identity provider.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "ontap_role": {
    "name": "admin"
  },
  "timestamp": "string"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "ontap_role": {
    "name": "admin"
  },
  "timestamp": "string"
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 5636243
| Provided ONTAP role is not configured.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role
```

```
ONTAP role to map with the external identity role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Role name


|===


[#security_external_role_mapping]
[.api-collapsible-fifth-title]
security_external_role_mapping

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this external-role-mapping entry.


|ontap_role
|link:#ontap_role[ontap_role]
a|ONTAP role to map with the external identity role.

* Introduced in: 9.16


|timestamp
|string
a|Date and time indicating when this external-role-mapping entry was
created.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage Google Cloud KMS

:leveloffset: +1


[[IDb00ccd5c09b56d5616dc209d9c716be9]]
= Manage Google Cloud KMS



== Overview

Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This feature
allows ONTAP to securely protect its encryption keys using Google Cloud
KMS.
In order to use Google Cloud KMS with ONTAP, a user must first deploy a
Google Cloud application with appropriate access to the Google Cloud KMS
and then provide
ONTAP with the necessary details, such as, project ID, key ring name,
location, key name and application credentials to allow ONTAP to
communicate
with the deployed Google Cloud application.
The properties `state`, `google_reachability` and `ekmip_reachability` are
considered advanced properties and are populated only when explicitly
requested.

== Examples
```

```
=== Enabling GCKMS for an SVM

The following example shows how to enable GCKMS at the SVM-scope. Note the
_return_records=true_ query parameter is used to obtain the newly created
key manager configuration.

----

# The API:
POST /api/security/gcp-kms

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/gcp-kms?return_records=true'
-H 'accept: application/hal+json' -d '{"svm":{"uuid":"f36ff553-e713-11ea-
bd56-005056bb4222" }, "project_id": "testProj",
"key_ring_name":"testKeyRing", "key_ring_location": "global", "key_name":
"key1", "application_credentials": "{\"client_email\":
\"my@account.email.com\", \"private_key\": \"ValidPrivateKey\"}"}'

# The response:
{
"num_records": 1,
"records": [
    {
    "uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
    "svm": {
        "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
        "name": "vs0"
    },
    "project_id": "testProj",
    "key_ring_name": "testKeyRing",
    "key_ring_location": "global",
    "key_name": "key1",
    "_links": {
        "self": {
        "href": "/api/security/gcp-kms/f72098a2-e908-11ea-bd56-
005056bb4222"
            }
        }
    }
]
}
----

'''
```

```
=== Retrieving all GCKMS configurations

The following example shows how to retrieve all GCKMS configurations.

----

# The API:
GET /api/security/gcp-kms

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/gcp-kms?fields=*'

# The response:
{
"records": [
    {
    "uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
    "scope": "svm",
    "svm": {
        "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
        "name": "vs0"
    },
    "project_id": "testProj",
    "key_ring_name": "testKeyRing",
    "key_ring_location": "global",
    "key_name": "key1",
    "_links": {
        "self": {
        "href": "/api/security/gcp-kms/f72098a2-e908-11ea-bd56-
005056bb4222"
        }
    }
    }
],
"num_records": 1,
"_links": {
    "self": {
    "href": "/api/security/gcp-kms?fields=*"
    }
}
}
----

'''

=== Retrieving a specific GCKMS configuration
```

```
The following example shows how to retrieve information for a specific
GCKMS configuration.

----

# The API:
GET /api/security/gcp-kms/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222?fields=*'

# The response:
{
"uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
"scope": "svm",
"svm": {
    "uuid": "f36ff553-e713-11ea-bd56-005056bb4222",
    "name": "vs0"
},
"project_id": "testProj",
"key_ring_name": "testKeyRing",
"key_ring_location": "global",
"key_name": "key1",
"_links": {
    "self": {
    "href": "/api/security/gcp-kms/f72098a2-e908-11ea-bd56-005056bb4222"
    }
}
}
----

'''

=== Retrieving a specific GCKMS's advanced properties

The following example shows how to retrieve advanced properties for a
specific GCKMS configuration.

----

# The API:
GET /api/security/gcp-kms/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222?fields=state,google_reachability,ekmip_reachability'
```

```
# The response:
{
"uuid": "f72098a2-e908-11ea-bd56-005056bb4222",
"state": {
    "cluster_state": false,
    "message": "The Google Cloud Key Management Service key protection is
unavailable on the following nodes: cluster1-node1.",
    "code": "65537708"
},
"google_reachability": {
    "reachable": true,
    "message": "",
    "code": "0"
},
"ekmip_reachability": [
    {
    "node": {
        "uuid": "d208115f-7721-11eb-bf83-005056bb150e",
        "name": "node1",
        "_links": {
            "self": {
            "href": "/api/cluster/nodes/d208115f-7721-11eb-bf83-
005056bb150e"
            }
        }
    },
    "reachable": true,
    "message": "",
    "code": "0"
    },
    {
    "node": {
        "uuid": "e208115f-7721-11eb-bf83-005056bb150e",
        "name": "node2",
        "_links": {
            "self": {
            "href": "/api/cluster/nodes/e208115f-7721-11eb-bf83-
005056bb150e"
            }
        }
    },
    "reachable": true,
    "message": "",
    "code": "0"
    }
```

```
],
"_links": {
    "self": {
    "href": "/api/security/gcp-kms/f72098a2-e908-11ea-bd56-005056bb4222"
    }
}
}
----

'''

=== Updating the application credentials of a specific GCKMS configuration

The following example shows how to update the application credentials for
a specific GCKMS configuration.

----

# The API:
PATCH /api/security/gcp-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222/' -d '{"application_credentials": "{\"client_email\":
\"new@account.com\", \"private_key\": \"ValidPrivateKey\"}"}'
----

'''

=== Updating the application credentials and applying a privileged account
for impersonation.

The following example shows how to set a privileged account on an existing
GCKMS configuration.

----

# The API:
PATCH /api/security/gcp-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222/' -d '{"application_credentials": "{\"client_email\":
\"unprivileged@account.com\", \"private_key\":
\"ValidPrivateKeyforUnprivilegedAccount\"}", "privileged_account":
"privileged@account.com"}'
----
```

```
'''

=== Updating the authentication method to use GCE metadata server for
retrieving short lived authentication tokens.

The following example shows how to change the authentication method on an
existing GCKMS configuration.

----

# The API:
PATCH /api/security/gcp-kms/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222/' -d '{"authentication_method":
"sa_credentials_attachment"}'
----

'''

=== Changing the GCKMS key

The following example shows how to change the GCKMS key.

----

# The API:
POST /api/security/gcp-kms/{uuid}/rekey-external

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222/rekey-external' -d '{"key_name": "key2",
"key_ring_location": "global", "key_ring_name": "testKeyRing2",
"project_id": "testProj2"}'
----

'''

=== Deleting a specific GCKMS configuration

The following example shows how to delete a specific GCKMS configuration.

----

# The API:
```

```
DELETE /api/security/gcp-kms/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/gcp-kms/f72098a2-e908-11ea-
bd56-005056bb4222'
----

'''

=== Restoring keys from a KMIP server

The following example shows how to restore keys for a GCKMS configuration.

----

# The API:
POST /api/security/gcp-kms/{uuid}/restore

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/gcp-kms/33820b57-ec90-11ea-
875e-005056bbf3f0/restore'
----

'''





[[ID07881e1bddbfe4cae167f356a122e2d2]]
= Retrieve a Google Cloud KMS configurations for all clusters and SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/gcp-kms`#

*Introduced In:* 9.9

Retrieves Google Cloud KMS configurations for all clusters and SVMs.

== Related ONTAP commands

* `security key-manager external gcp show`
* `security key-manager external gcp check`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|port
|integer
|query
|False
a|Filter by port

* Introduced in: 9.14


|proxy_type
|string
|query
|False
a|Filter by proxy_type


|project_id
|string
|query
|False
a|Filter by project_id


|uuid
|string
|query
|False
a|Filter by uuid


|ekmip_reachability.message
|string
|query
|False
a|Filter by ekmip_reachability.message


|ekmip_reachability.node.name
```

```
|string
|query
|False
a|Filter by ekmip_reachability.node.name


|ekmip_reachability.node.uuid
|string
|query
|False
a|Filter by ekmip_reachability.node.uuid


|ekmip_reachability.reachable
|boolean
|query
|False
a|Filter by ekmip_reachability.reachable


|ekmip_reachability.code
|string
|query
|False
a|Filter by ekmip_reachability.code


|key_name
|string
|query
|False
a|Filter by key_name


|authentication_method
|string
|query
|False
a|Filter by authentication_method

* Introduced in: 9.17


|state.code
|string
|query
|False
```

```
a|Filter by state.code


|state.cluster_state
|boolean
|query
|False
a|Filter by state.cluster_state


|state.message
|string
|query
|False
a|Filter by state.message


|key_ring_location
|string
|query
|False
a|Filter by key_ring_location


|verify_ip
|boolean
|query
|False
a|Filter by verify_ip

* Introduced in: 9.14


|proxy_port
|integer
|query
|False
a|Filter by proxy_port


|key_ring_name
|string
|query
|False
a|Filter by key_ring_name
```

```
|gce_metadata_server
|string
|query
|False
a|Filter by gce_metadata_server

* Introduced in: 9.17


|google_reachability.reachable
|boolean
|query
|False
a|Filter by google_reachability.reachable


|google_reachability.code
|string
|query
|False
a|Filter by google_reachability.code


|google_reachability.message
|string
|query
|False
a|Filter by google_reachability.message


|proxy_username
|string
|query
|False
a|Filter by proxy_username


|cloudkms_host
|string
|query
|False
a|Filter by cloudkms_host

* Introduced in: 9.14


|privileged_account
```

```
|string
|query
|False
a|Filter by privileged_account

* Introduced in: 9.14


|scope
|string
|query
|False
a|Filter by scope


|proxy_host
|string
|query
|False
a|Filter by proxy_host


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|oauth_url
|string
|query
|False
a|Filter by oauth_url

* Introduced in: 9.14


|oauth_host
|string
```

```
|query
|False
a|Filter by oauth_host

* Introduced in: 9.14


|verify_host
|boolean
|query
|False
a|Filter by verify_host

* Introduced in: 9.14


|caller_account
|string
|query
|False
a|Filter by caller_account

* Introduced in: 9.14


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#gcp_kms[gcp_kms]]
```

```
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_method": "application_credentials_key",
      "caller_account": "myaccount@myproject.com",
      "cloudkms_host": "cloudkms.googleapis.com",
      "ekmip_reachability": [
        {
          "code": "346758",
          "message": "embedded KMIP server status unavailable on node.",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "gce_metadata_server":
"http://169.254.169.254/computeMetadata/v1/instance/service-
```

```
accounts/default/token",
      "google_reachability": {
        "code": "346758",
        "message": "Google Cloud KMS is not reachable from all nodes -
<reason>."
      },
      "key_name": "cryptokey1",
      "key_ring_location": "global",
      "key_ring_name": "gcpapp1-keyring",
      "oauth_host": "oauth2.googleapis.com",
      "oauth_url": "https://oauth2.googleapis.com/token",
      "port": 443,
      "privileged_account":
"myserviceaccount@myproject.iam.gserviceaccount.com",
      "project_id": "gcpapp1",
      "proxy_host": "proxy.eng.com",
      "proxy_port": 1234,
      "proxy_type": "http",
      "proxy_username": "proxyuser",
      "scope": "string",
      "state": {
        "code": "346758",
        "message": "Top-level internal key protection key (KEK) is
unavailable on the following nodes with the associated reasons: Node:
node1. Reason: No volumes created yet for the SVM. Wrapped KEK status will
be available after creating encrypted volumes."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.

| 65537730
| The Google Cloud Key Management Service is unreachable from one or more
nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
```

```
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#google_reachability]
```

```
[.api-collapsible-fifth-title]
google_reachability

Indicates whether or not the Google Cloud KMS is reachable from all nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Google Cloud KMS
is reachable from all nodes in the cluster.


|message
|string
a|Set to the error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Google Cloud KMS is reachable from all nodes of the
cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
```

```
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when Google Cloud KMS key protection is available on all
nodes of the cluster.


|code
|string
a|Error code corresponding to the status message. Returns 0 if Google
Cloud KMS key protection is available in all nodes of the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#gcp_kms]
[.api-collapsible-fifth-title]
gcp_kms

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Google Cloud KMS authentication method.


|caller_account
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
```

```
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.


|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|key_name
|string
a|Key Identifier of Google Cloud KMS key encryption key.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name of the deployed Google Cloud application.


|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.
```

```
|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|project_id
|string
a|Google Cloud project (application) ID of the deployed Google Cloud
application that has appropriate access to the Google Cloud KMS.


|proxy_host
|string
a|Proxy host name.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
```

```
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
a|Verify the identity of the Google Cloud KMS IP address.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb2ed1e56349fbadb74d37b6db203e90d]]
= Create a Google Cloud KMS configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/gcp-kms`#

*Introduced In:* 9.9

Configures the Google Cloud KMS configuration for the specified SVM.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create a Google
Cloud KMS.
* `project_id` - Google Cloud project (application) ID of the deployed
Google Cloud application with appropriate access to the Google Cloud KMS.
* `key_ring_name` - Google Cloud KMS key ring name of the deployed Google
Cloud application with appropriate access to the specified Google Cloud
KMS.
* `key_ring_location` - Google Cloud KMS key ring location.
* `key_name`- Key Identifier of the Google Cloud KMS key encryption key.
* `application_credentials` - The Google Cloud application's service
account credentials required to access the specified KMS. This is a JSON
file containing an email address and the private key of the service
account holder. The field is required if the `authentication_method` is
set to `application_credentials_key`.

== Optional properties

* `proxy_type` - Type of proxy (http/https) if proxy configuration is
used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `port` - Authorization server and Google Cloud KMS port number.
* `cloudkms_host` - Google Cloud KMS host subdomain.
* `oauth_host` - Open authorization server host name.
* `oauth_url` - Open authorization URL for the access token.
* `privileged_account` - Account used to impersonate Google Cloud KMS
requests.
* `verify_ip` - Verify identity of Google Cloud KMS IP address.
* `verify_host` - Verify the identity of the Google Cloud KMS host name.
* `authentication_method` - The authentication method to use for
retrieving short lived authentication tokens.
* `gce_metadata_server` - A custom metadata server URL used for retrieving
short lived authentication tokens if the default service account is not
used. This is only applicable when the `authentication_method` is set to
```

```
`sa_credentials_attachment`.

== Related ONTAP commands

* `security key-manager external gcp enable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_credentials
|string
a|The Google Cloud application's service account credentials required to
access the specified KMS. The client_email and private_key fields of the
service account holder are required. The credentials are required if the
`authentication_method` is set to `application_credentials_key`.


|authentication_method
|string
a|Google Cloud KMS authentication method.


|caller_account
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.
```

```
|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|key_name
|string
a|Key Identifier of Google Cloud KMS key encryption key.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name of the deployed Google Cloud application.


|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.


|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|project_id
```

```
|string
a|Google Cloud project (application) ID of the deployed Google Cloud
application that has appropriate access to the Google Cloud KMS.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
a|Verify the identity of the Google Cloud KMS IP address.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "application_credentials":
"{\"private_key\":\"ValidPrivateKey\",\"client_email\":\"my@account.email.
com\"}",
  "authentication_method": "application_credentials_key",
  "caller_account": "myaccount@myproject.com",
  "cloudkms_host": "cloudkms.googleapis.com",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "gce_metadata_server":
```

```
"http://169.254.169.254/computeMetadata/v1/instance/service-
accounts/default/token",
  "google_reachability": {
    "code": "346758",
    "message": "Google Cloud KMS is not reachable from all nodes -
<reason>."
  },
  "key_name": "cryptokey1",
  "key_ring_location": "global",
  "key_ring_name": "gcpapp1-keyring",
  "oauth_host": "oauth2.googleapis.com",
  "oauth_url": "https://oauth2.googleapis.com/token",
  "port": 443,
  "privileged_account":
"myserviceaccount@myproject.iam.gserviceaccount.com",
  "project_id": "gcpapp1",
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "state": {
    "code": "346758",
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```


Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537703
| The Google Cloud Key Management Service is not supported for the admin
Vserver.

| 65537704
| The Google Cloud Key Management Service is not supported in MetroCluster
configurations.

| 65537706
| Internal error. Failed to the encrypt the application credentials.

| 65537713
| Internal Error. Failed to store the application credentials.

| 65537719
| Failed to enable the Google Cloud Key Management Service for SVM
+++<svm-name>+++because invalid application credentials were
provided.+++</svm-name>+++

| 65537720
| Failed to configure Google Cloud Key Management Service for SVM +++<svm-
name>+++because a key manager has already been configured for this SVM.
Use the REST API GET method \"/api/security/gcp-kms\" to view all of the
configured key managers.+++</svm-name>+++

| 65537740
| The privileged account must be an email address or an empty string.

| 65537749
| The application credentials field must be in valid JSON format.

| 65537753
| The "key_name" value contains invalid characters. A forward slash '/' is
not allowed in the key name.

| 65539704
```

```
| The key manager cannot be configured because the SVM has NAE volumes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
```

```
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
```

```
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#google_reachability]
[.api-collapsible-fifth-title]
google_reachability

Indicates whether or not the Google Cloud KMS is reachable from all nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Google Cloud KMS
is reachable from all nodes in the cluster.


|message
|string
a|Set to the error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Google Cloud KMS is reachable from all nodes of the
cluster.


|===


[#state]
```

```
[.api-collapsible-fifth-title]
state

Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when Google Cloud KMS key protection is available on all
nodes of the cluster.


|code
|string
a|Error code corresponding to the status message. Returns 0 if Google
Cloud KMS key protection is available in all nodes of the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#gcp_kms]
[.api-collapsible-fifth-title]
gcp_kms

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_credentials
|string
a|The Google Cloud application's service account credentials required to
access the specified KMS. The client_email and private_key fields of the
service account holder are required. The credentials are required if the
`authentication_method` is set to `application_credentials_key`.


|authentication_method
|string
a|Google Cloud KMS authentication method.
```

```
|caller_account
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.


|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|key_name
|string
a|Key Identifier of Google Cloud KMS key encryption key.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name of the deployed Google Cloud application.
```

```
|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.


|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|project_id
|string
a|Google Cloud project (application) ID of the deployed Google Cloud
application that has appropriate access to the Google Cloud KMS.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.
```

```
|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
a|Verify the identity of the Google Cloud KMS IP address.


|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID44448f480c4a1b6ea0aa1399d2bb781c]]
= Re-key the external key in the key hierarchy for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/gcp-kms/{gcp_kms.uuid}/rekey-external`#

*Introduced In:* 9.11

Rekeys the external key in the key hierarchy for an SVM with a Google
Cloud KMS configuration.

== Required properties
```

```
* `key_name` - Google Cloud KMS key name.

== Optional properties

* `project_id` - Google Cloud KMS project ID.
* `key_ring_location` - Google Cloud KMS key ring location.
* `key_ring_name` - Google Cloud KMS key ring name.

== Related ONTAP commands

* `security key-manager external gcp rekey-external`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|gcp_kms.uuid
|string
|path
|True
a|UUID of the existing Google Cloud KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0
```

```
|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_name
|string
a|Google Cloud KMS key name.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name.


|project_id
|string
a|Google Cloud KMS project ID.


|===


.Example request
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "key_name": "cryptokey1",
  "key_ring_location": "location1",
  "key_ring_name": "keyring1",
  "project_id": "project1"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537556
| ONTAP is not able to successfully encrypt or decrypt because the
configured external key manager for this SVM is in a blocked state.
Possible reasons for a blocked state include the top-level external key
protection key not found, disabled or having insufficient privileges.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65537721
| Google Cloud KMS is not configured for the given SVM.

| 65537729
| External rekey failed on one or more nodes. Use the REST API POST method
"/api/security/gcp-kms/\{uuid}/rekey-external" to try the rekey operation
again.

| 65537753
| The "key_name" value contains invalid characters. A forward slash '/' is
not allowed in the key name.

| 65537755
| Google Cloud KMS configuration details were not found for the SVM.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
```

```
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#gcp_kms_key]
[.api-collapsible-fifth-title]
gcp_kms_key

[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_name
|string
a|Google Cloud KMS key name.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name.


|project_id
|string
a|Google Cloud KMS project ID.


|===
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID6f91ef126cc42beb1635382d432afb53]]
= Delete a Google Cloud KMS configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/gcp-kms/{uuid}`#

*Introduced In:* 9.9

Deletes a Google Cloud KMS configuration.

== Related ONTAP commands

* `security key-manager external gcp disable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Google Cloud KMS UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
```

```
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536817
| Internal error. Failed to determine if it is safe to disable key
manager.

| 65536827
| Internal error. Failed to determine if the given SVM has any encrypted
volumes.

| 65536834
| Internal error. Failed to get existing key-server details for the given
SVM.

| 65536867
| Volume encryption keys (VEK) for one or more encrypted volumes are
stored on the key manager configured for the given SVM.
```

```
| 65536883
| Internal error. Volume encryption key is missing for a volume.

| 65536884
| Internal error. Volume encryption key is invalid for a volume.

| 65537721
| The Google Cloud Key Management Service is not configured for the SVM.

| 196608080
| One or more nodes in the cluster have the root volume encrypted using
NVE (NetApp Volume Encryption).

| 196608301
| Internal error. Failed to get encryption type.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
```

```
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDc3ae4a60d6a228c8090d7dc7600a984c]]
= Retrieve the Google Cloud KMS configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/gcp-kms/{uuid}`#

*Introduced In:* 9.9

Retrieves the Google Cloud KMS configuration for the SVM specified by the
UUID.

== Related ONTAP commands

* `security key-manager external gcp show`
* `security key-manager external gcp check`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Google Cloud KMS UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_method
|string
a|Google Cloud KMS authentication method.


|caller_account
```

```
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.


|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|key_name
|string
a|Key Identifier of Google Cloud KMS key encryption key.


|key_ring_location
|string
a|Google Cloud KMS key ring location.


|key_ring_name
|string
a|Google Cloud KMS key ring name of the deployed Google Cloud application.
```

```
|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.


|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|project_id
|string
a|Google Cloud project (application) ID of the deployed Google Cloud
application that has appropriate access to the Google Cloud KMS.


|proxy_host
|string
a|Proxy host name.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
```

```
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
a|Verify the identity of the Google Cloud KMS IP address.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_method": "application_credentials_key",
  "caller_account": "myaccount@myproject.com",
  "cloudkms_host": "cloudkms.googleapis.com",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "gce_metadata_server":
"http://169.254.169.254/computeMetadata/v1/instance/service-
accounts/default/token",
  "google_reachability": {
    "code": "346758",
    "message": "Google Cloud KMS is not reachable from all nodes -
<reason>."
  },
  "key_name": "cryptokey1",
  "key_ring_location": "global",
  "key_ring_name": "gcpapp1-keyring",
  "oauth_host": "oauth2.googleapis.com",
  "oauth_url": "https://oauth2.googleapis.com/token",
  "port": 443,
  "privileged_account":
"myserviceaccount@myproject.iam.gserviceaccount.com",
  "project_id": "gcpapp1",
  "proxy_host": "proxy.eng.com",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "state": {
    "code": "346758",
```

```
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537551
| Top-level internal key protection key (KEK) unavailable on one or more
nodes.

| 65537552
| Embedded KMIP server status not available.

| 65537730
| The Google Cloud Key Management Service is unreachable from one or more
nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===
```

```
[#google_reachability]
[.api-collapsible-fifth-title]
google_reachability

Indicates whether or not the Google Cloud KMS is reachable from all nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Google Cloud KMS
is reachable from all nodes in the cluster.


|message
|string
a|Set to the error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Google Cloud KMS is reachable from all nodes of the
cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
```

```
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when Google Cloud KMS key protection is available on all
nodes of the cluster.


|code
|string
a|Error code corresponding to the status message. Returns 0 if Google
Cloud KMS key protection is available in all nodes of the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID9058f544f8344a844c7ee70a562104c6]]
= Update the Google Cloud KMS configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/gcp-kms/{uuid}`#

*Introduced In:* 9.9

Updates the Google Cloud KMS configuration.

== Optional properties

* `application_credentials` - New credentials used to verify the
```

```
application's identity to the Google Cloud KMS.
* `proxy_type` - Type of proxy (http/https) if proxy configuration is
used.
* `proxy_host` - Proxy hostname if proxy configuration is used.
* `proxy_port` - Proxy port number if proxy configuration is used.
* `port` - Authorization server and Google Cloud KMS port number.
* `proxy_username` - Proxy username if proxy configuration is used.
* `proxy_password` - Proxy password if proxy configuration is used.
* `project_id` - Google Cloud project (application) ID of the deployed
Google Cloud application with appropriate access to the Google Cloud KMS.
* `key_ring_name` - Google Cloud KMS key ring name of the deployed Google
Cloud application with appropriate access to the specified Google Cloud
KMS.
* `key_ring_location` - Google Cloud KMS key ring location.
* `cloudkms_host` - Google Cloud KMS host subdomain.
* `oauth_host` - Open authorization server host name.
* `oauth_url` - Open authorization URL for the access token.
* `verify_host` - Verify the identity of the Google Cloud KMS host name.
* `verify_ip ` - Verify identity of Google Cloud KMS IP address.
* `privileged_account` - Account used to impersonate Google Cloud KMS
requests.
* `authentication_method` - The authentication method to use for
retrieving short lived authentication tokens.
* `gce_metadata_server` - A custom metadata server URL used for retrieving
short lived authentication tokens if the default service account is not
used. This is only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.

== Related ONTAP commands

* `security key-manager external gcp update-credentials`
* `security key-manager external gcp update-config`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
```

```
|string
|path
|True
a|Google Cloud KMS UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_credentials
|string
a|The Google Cloud application's service account credentials required to
access the specified KMS. The client_email and private_key fields of the
service account holder are required. The credentials are required if the
`authentication_method` is set to `application_credentials_key`.


|authentication_method
```

```
|string
a|Google Cloud KMS authentication method.


|caller_account
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.


|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.
```

```
|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.


|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
a|Verify the identity of the Google Cloud KMS IP address.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "application_credentials":
"{\"private_key\":\"ValidPrivateKey\",\"client_email\":\"my@account.email.
com\"}",
  "authentication_method": "application_credentials_key",
  "caller_account": "myaccount@myproject.com",
  "cloudkms_host": "cloudkms.googleapis.com",
  "ekmip_reachability": [
    {
      "code": "346758",
      "message": "embedded KMIP server status unavailable on node.",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "gce_metadata_server":
```

```
"http://169.254.169.254/computeMetadata/v1/instance/service-
accounts/default/token",
  "google_reachability": {
    "code": "346758",
    "message": "Google Cloud KMS is not reachable from all nodes -
<reason>."
  },
  "oauth_host": "oauth2.googleapis.com",
  "oauth_url": "https://oauth2.googleapis.com/token",
  "port": 443,
  "privileged_account":
"myserviceaccount@myproject.iam.gserviceaccount.com",
  "proxy_host": "proxy.eng.com",
  "proxy_password": "proxypassword",
  "proxy_port": 1234,
  "proxy_type": "http",
  "proxy_username": "proxyuser",
  "scope": "string",
  "state": {
    "code": "346758",
    "message": "Top-level internal key protection key (KEK) is unavailable
on the following nodes with the associated reasons: Node: node1. Reason:
No volumes created yet for the SVM. Wrapped KEK status will be available
after creating encrypted volumes."
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```


Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 65537541
| No inputs were provided for the patch request.

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.

| 65537605
| Failed to establish connectivity with the cloud key management service.

| 65537713
| Internal Error. Failed to store the application credentials.

| 65537714
| The "application_credentials" field must be specified.

| 65537721
| The Google Cloud Key Management Service is not configured for the SVM.

| 65537724
| Failed to update the Google Cloud Key Management Service because invalid
application credentials were provided.

| 65537757
| The "key_name", "project_id", "key_ring_name" and "key_ring_location"
fields cannot be modified directly. Use the REST API POST method
"/api/security/gcp-kms/\{uuid}/rekey-external" to modify these fields.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
```

```
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#ekmip_reachability]
[.api-collapsible-fifth-title]
ekmip_reachability

Provides the connectivity status for the given SVM on the given node to
all EKMIP servers configured on all nodes of the cluster.
This is an advanced property; there is an added computational cost to
```

```
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if a given SVM is
able to communicate to the EKMIP servers of all of the nodes in the
cluster.


|message
|string
a|Error message set when cluster-wide EKMIP server availability from the
given SVM and node is false.


|node
|link:#node[node]
a|

|reachable
|boolean
a|Set to true if the given SVM on the given node is able to communicate to
all EKMIP servers configured on all nodes in the cluster.


|===


[#google_reachability]
[.api-collapsible-fifth-title]
google_reachability

Indicates whether or not the Google Cloud KMS is reachable from all nodes
in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
```

```
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Code corresponding to the error message. Returns a 0 if Google Cloud KMS
is reachable from all nodes in the cluster.


|message
|string
a|Set to the error message when 'reachable' is false.


|reachable
|boolean
a|Set to true if the Google Cloud KMS is reachable from all nodes of the
cluster.


|===


[#state]
[.api-collapsible-fifth-title]
state

Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|cluster_state
|boolean
a|Set to true when Google Cloud KMS key protection is available on all
nodes of the cluster.


|code
|string
a|Error code corresponding to the status message. Returns 0 if Google
Cloud KMS key protection is available in all nodes of the cluster.


|message
|string
a|Error message set when top-level internal key protection key (KEK)
availability on cluster is false.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#gcp_kms]
[.api-collapsible-fifth-title]
gcp_kms

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application_credentials
|string
a|The Google Cloud application's service account credentials required to
access the specified KMS. The client_email and private_key fields of the
service account holder are required. The credentials are required if the
`authentication_method` is set to `application_credentials_key`.


|authentication_method
|string
a|Google Cloud KMS authentication method.


|caller_account
|string
a|Google Cloud KMS caller account email


|cloudkms_host
|string
a|Google Cloud KMS host subdomain.


|ekmip_reachability
|array[link:#ekmip_reachability[ekmip_reachability]]
a|

|gce_metadata_server
```

```
|string
a|A custom metadata server URL used for retrieving short lived
authentication tokens if the default service account is not used. This is
only applicable when the `authentication_method` is set to
`sa_credentials_attachment`.


|google_reachability
|link:#google_reachability[google_reachability]
a|Indicates whether or not the Google Cloud KMS is reachable from all
nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|oauth_host
|string
a|Open authorization server host name.


|oauth_url
|string
a|Open authorization URL for the access token.


|port
|integer
a|Authorization server and Google Cloud KMS port number.


|privileged_account
|string
a|Google Cloud KMS account to impersonate.


|proxy_host
|string
a|Proxy host name.


|proxy_password
|string
a|Proxy password. Password is not audited.
```

```
|proxy_port
|integer
a|Proxy port number.


|proxy_type
|string
a|Type of proxy.


|proxy_username
|string
a|Proxy username.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|link:#state[state]
a|Google Cloud Key Management Services is a cloud key management service
(KMS) that provides a secure store for encryption keys. This object
indicates whether or not the Google Cloud KMS key protection is available
on all nodes in the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|uuid
|string
a|A unique identifier for the Google Cloud KMS.


|verify_host
|boolean
a|Verify the identity of the Google Cloud KMS host name.


|verify_ip
|boolean
```

```
a|Verify the identity of the Google Cloud KMS IP address.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID717ac4bc69491bba31a2bc080ec616bd]]
= Re-key the internal key in the key hierarchy for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/gcp-kms/{uuid}/rekey-internal`#

*Introduced In:* 9.10

Rekeys the internal key in the key hierarchy for an SVM with a Google
Cloud KMS configuration.

== Related ONTAP commands

* `security key-manager external gcp rekey-internal`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|UUID of the existing Google Cloud KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537547
| One or more volume encryption keys for encrypted volumes of this data
SVM are stored in the key manager configured for the admin SVM. Use the
REST API POST method to migrate this data SVM's keys from the admin SVM's
key manager to this data SVM's key manager before running the rekey
operation.
```

```
| 65537556
| ONTAP is not able to successfully encrypt or decrypt because the
configured external key manager for this SVM is in a blocked state.
Possible reasons for a blocked state include the top-level external key
protection key not found, disabled or having insufficient privileges.

| 65537559
| There are no existing internal keys for the SVM. A rekey operation is
allowed for an SVM with one or more encryption keys.

| 65537610
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent rekeying for the SVM. Use the REST API PATCH method
"/api/security/key-stores/\{uuid}" to re-run and complete the operation.

| 65537721
| Google Cloud KMS is not configured for the given SVM.

| 65539436
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being initialized. Wait until the keystore is in the
active state, and rerun the rekey operation.

| 65539437
| Rekey cannot be performed on the SVM while the enabled keystore
configuration is being disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID3cf4232c211235b36d2b75753fe29859]]
= Restore keys for an SVM from a Google Cloud KMS

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/gcp-kms/{uuid}/restore`#

*Introduced In:* 9.10

Restores the keys for an SVM from a configured Google Cloud KMS.

== Related ONTAP commands

* `security key-manager external gcp restore`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the existing Google Cloud KMS configuration.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
```

```
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537544
| Missing wrapped top-level internal key protection key (KEK) from
internal database.

| 65537721
| The Google Cloud Key Management Service is not configured for the given
SVM.

| 65537722
| Failed to restore keys on the following nodes.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage group-to-role mapping entries

:leveloffset: +1
```

```
[[ID85c033e842ad812e6d4b6e04f489f8f0]]
= Manage group-to-role mapping entries



=== Overview

This API is used to manage and map groups to ONTAP roles. The POST request
creates a group to role mapping. Various responses are shown in the
examples below.
 +

== Examples

=== Creating a group to role mapping entry

The following output shows how to create a group to role mapping entry.

----

# The API:
/api/security/group/role-mappings

# The call:
curl -X POST "https://<mgmt-ip>/api/security/group/role-mappings" -d '{
"group_id": 1, "ontap_role": { "name": "admin" }, "comment": "Group1 admin
role"}'
----

=== Retrieving all group to role mapping entries

The following output shows all group to role mapping entries.

----

# The API:
/api/security/group/role-mappings

# The call:
curl -X GET "https://<mgmt-ip>/api/security/group/role-mappings"

# The response:
{
 "records": [
 {
```

```
    "group_id": 1,
    "ontap_role": {
      "name": "admin"
    }
  }
  ],
  "num_records": 1
}
```


```
[[ID9dd6c443681d75a68aedd381169a312d]]
= Retrieve all group-to-role mapping entries

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/group/role-mappings`#

*Introduced In:* 9.16

Retrieves all the group to role mapping entries.

== Related ONTAP commands

* `security login group role-mapping show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|ontap_role.name
|string
|query
|False
a|Filter by ontap_role.name
```

```
|comment
|string
|query
|False
a|Filter by comment


|group_id
|integer
|query
|False
a|Filter by group_id

* Max value: 4294967295
* Min value: 1


|scope
|string
|query
|False
a|Filter by scope


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#group_role_mappings[group_role_mappings]]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "group_id": 1,
      "ontap_role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "scope": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role

Role name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Role name


|===


[#group_role_mappings]
[.api-collapsible-fifth-title]
group_role_mappings

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|Any comment regarding this group entry.


|group_id
|integer
```

```
a|Group ID.


|ontap_role
|link:#ontap_role[ontap_role]
a|Role name.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID2589ee0ecc2831e15cf4f1bfb5bcea02]]
= Create a group-to-role mapping entry

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/group/role-mappings`#

*Introduced In:* 9.16

Creates a group to role mapping entry.

== Required properties

* `group_id`
* `ontap_role`

== Optional properties
```

```
* `comment`

== Related ONTAP commands

* `security login group role-mapping create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|group_id
|integer
a|Group ID.
```

```
|ontap_role
|link:#ontap_role[ontap_role]
a|Role name.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "group_id": 1,
  "ontap_role": {
    "name": "admin"
  },
  "scope": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
```

```
|array[link:#group_role_mappings[group_role_mappings]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "comment": "string",
      "group_id": 1,
      "ontap_role": {
        "name": "admin"
      },
      "scope": "string"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5636242
| The configured group ID does not exist in the group configuration.

| 5636243
| The configured role name does not exist in the SVM where the group ID is
configured.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#ontap_role]
[.api-collapsible-fifth-title]
```

```
ontap_role

Role name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Role name


|===


[#group_role_mappings]
[.api-collapsible-fifth-title]
group_role_mappings

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|group_id
|integer
a|Group ID.


|ontap_role
|link:#ontap_role[ontap_role]
a|Role name.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
```

```
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID1dab926fef2d1a62118819fac088b15c]]
= Delete a group-to-role mapping entry

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/group/role-mappings/{group_id}/{ontap_role.name}`#

*Introduced In:* 9.16

Deletes a group to role mapping entry.

== Related ONTAP commands

* `security login group role-mapping delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|group_id
```

```
|integer
|path
|True
a|Group ID.

* format: int64
* Max value: 4294967295
* Min value: 1


|ontap_role.name
|string
|path
|True
a|Role name.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd66df4fb21530938af260a81a8de1d68]]
= Retrieve a group-to-role mapping entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/group/role-mappings/{group_id}/{ontap_role.name}`#

*Introduced In:* 9.16

Retrieves a group to role mapping entry.

== Related ONTAP commands

* `security login group role-mapping show`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|group_id
|integer
|path
|True
a|Group ID.

* format: int64
* Max value: 4294967295
* Min value: 1


|ontap_role.name
|string
|path
|True
a|Role name.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|comment
|string
a|Any comment regarding this group entry.


|group_id
|integer
a|Group ID.


|ontap_role
|link:#ontap_role[ontap_role]
a|Role name.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "group_id": 1,
  "ontap_role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
```

```
  "scope": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role

Role name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Role name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID94f9dc176c0fee147741090092f02292]]
= Update a configured group role-mapping entry

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/group/role-mappings/{group_id}/{ontap_role.name}`#

*Introduced In:* 9.16

Updates a configured group role-mapping.

== Related ONTAP commands

* `security login group role-mapping modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|group_id
|integer
|path
|True
a|Group ID.

* format: int64
* Max value: 4294967295
* Min value: 1


|ontap_role.name
|string
|path
|True
a|Role name.

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "scope": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#ontap_role]
[.api-collapsible-fifth-title]
ontap_role

Role name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Role name


|===


[#group_role_mappings]
[.api-collapsible-fifth-title]
group_role_mappings
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage group configurations

:leveloffset: +1


[[ID91d7bd17f4804f6fa20ace391f9e0a55]]
= Manage group configurations
```

```
=== Overview

This API is used to manage group configurations in ONTAP. The POST request
creates a group configuration. Various responses are shown in the examples
below.
 +

== Examples

=== Creating a group configuration

The following output shows how to create a group configuration.

----

# The API:
/api/security/groups

# The call:
curl -X POST "https://<mgmt-ip>/api/security/groups"  -d '{ "name":
"AzureGroup1", "owner": { "name": "C1_sti213-vsim-sr023a_1718680001",
"uuid": "116127b1-2d21-11ef-a5e1-005056ae1bc2" }, "type": "entra", "uuid":
"3fa85f64-5717-4562-b3fc-2c963f66afa6", "comment": "Azure Group"}'
----

=== Retrieving all group configurations

The following output shows all group configurations.

----

# The API:
/api/security/groups

# The call:
curl -X GET "https://<mgmt-ip>/api/security/groups"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "116127b1-2d21-11ef-a5e1-005056ae1bc2",
      "name": "C1_sti213-vsim-sr023a_1718680001"
    },
    "name": "AzureGroup1",
```

```
    "type": "entra",
    "id": 1,
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
],
"num_records": 1
}
----
```


```
[[ID35fb7afcf152fefd0e7f3bf67974ab92]]
= Retrieve all group entries

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/groups`#

*Introduced In:* 9.16

Retrieves all group entries.

== Related ONTAP commands

* `security login group show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|type
|string
|query
|False
a|Filter by type


|id
```

```
|integer
|query
|False
a|Filter by id


|name
|string
|query
|False
a|Filter by name

* maxLength: 64
* minLength: 1


|scope
|string
|query
|False
a|Filter by scope


|uuid
|string
|query
|False
a|Filter by uuid


|comment
|string
|query
|False
a|Filter by comment


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
```

```
|False
a|Filter by owner.uuid


|create_time
|string
|query
|False
a|Filter by create_time


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_group[security_group]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "Azure Group",
      "create_time": "string",
      "id": 0,
      "name": "AzureGroup1",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "string",
      "type": "entra",
      "uuid": "string"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Group owner. Used to identify a cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_group]
[.api-collapsible-fifth-title]
security_group

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.


|name
|string
a|Group name.


|owner
|link:#owner[owner]
a|Group owner. Used to identify a cluster or an SVM.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|type
|string
a|Group type.


|uuid
|string
a|Group UUID from external Active Directory.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID41e75a5376237f695b1ee0f7f11224b9]]
= Create a group entry

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/groups`#

*Introduced In:* 9.16

Creates a group entry.

== Required properties

* `name`

== Optional properties

* `owner`
* `type`
* `uuid`
* `comment`

== Related ONTAP commands

* `security login group create`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.


|name
|string
```

```
a|Group name.


|owner
|link:#owner[owner]
a|Group owner. Used to identify a cluster or an SVM.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|type
|string
a|Group type.


|uuid
|string
a|Group UUID from external Active Directory.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "Azure Group",
  "create_time": "string",
  "id": 0,
  "name": "AzureGroup1",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "type": "entra",
  "uuid": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#security_group[security_group]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "comment": "Azure Group",
      "create_time": "string",
      "id": 0,
      "name": "AzureGroup1",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "string",
      "type": "entra",
      "uuid": "string"
    }
  ]
}
====

=== Headers
```

```
[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5636231
| Failed to create the group because the maximum number of allowed group
IDs has been reached.

| 5636232
| The group name does not satisfy the length requirement.

| 5636235
| Duplicate UUID.

| 5636244
| Failed to delete the group because it is part of the group role mapping
configuration.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

Group owner. Used to identify a cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_group]
[.api-collapsible-fifth-title]
security_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.


|name
|string
a|Group name.


|owner
|link:#owner[owner]
a|Group owner. Used to identify a cluster or an SVM.


|scope
|string
```

```
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|type
|string
a|Group type.


|uuid
|string
a|Group UUID from external Active Directory.


|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID745ac56347d255f87f86fea82cf3cb29]]
= Delete a group entry

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/groups/{owner.uuid}/{name}/{type}`#

*Introduced In:* 9.16

Deletes a group entry.

== Related ONTAP commands

* `security login group delete`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|name
|string
|path
|True
a|

|type
|string
|path
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description
```

```
| 5636244
| The group is part of the role mapping configuration. Delete the mapping
first and then delete the group.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6dea5adb4bab270e2cff345ac054d231]]
= Retrieve a group entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/groups/{owner.uuid}/{name}/{type}`#

*Introduced In:* 9.16

Retrieves a group entry.

== Related ONTAP commands

* `security login group show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
```

```
a|Group owner. Used to identify a cluster or an SVM.


|name
|string
|path
|True
a|Group name.


|type
|string
|path
|True
a|Group type.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
```

```
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.


|name
|string
a|Group name.


|owner
|link:#owner[owner]
a|Group owner. Used to identify a cluster or an SVM.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|type
|string
a|Group type.


|uuid
|string
a|Group UUID from external Active Directory.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
```

```
  },
  "comment": "Azure Group",
  "create_time": "string",
  "id": 0,
  "name": "AzureGroup1",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "type": "entra",
  "uuid": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
```

```
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Group owner. Used to identify a cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4f0d99879644ecde8f8c2e2beae2fffc]]
= Update configured group information

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/groups/{owner.uuid}/{name}/{type}`#

*Introduced In:* 9.16

Updates configured group information.

== Related ONTAP commands

* `security login group modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|owner.uuid
|string
|path
|True
a|Group owner. Used to identify a cluster or an SVM.


|name
|string
|path
|True
a|Group name.


|type
|string
|path
|True
a|Group type.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.
```

```
|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|uuid
|string
a|Group UUID from external Active Directory.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "Azure Group",
  "create_time": "string",
  "id": 0,
  "scope": "string",
  "uuid": "string"
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5636235
| Duplicate UUID.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

Group owner. Used to identify a cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_group]
[.api-collapsible-fifth-title]
security_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Any comment regarding this group entry.


|create_time
|string
a|Date and time indicating when this group entry was created.


|id
|integer
a|Group ID which is unique per group across the system. It is a read-only
field and it automatically assigns the next available unique number.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.
```

```
|uuid
|string
a|Group UUID from external Active Directory.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage HA network security

:leveloffset: +1


[[ID8b317c002cd6507ae6690e0bfe15f52e]]
= Manage HA network security



== Overview

You can use the security ha-network API endpoint to modify the HA network
security configuration for NVLog traffic.
The following operations are supported:

* GET to retrieve the HA network security status: GET security/ha-network
* PATCH to update the HA network security configuration: PATCH
security/ha-network
```

```
[[ID3536922f7b4bbfae67a0b180603d7de3]]
= Retrieve the HA network security configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ha-network`#

*Introduced In:* 9.18

Retrieves the HA network security configuration.

== Related ONTAP commands

* 'security ha-network show'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|enabled
|boolean
a|Indicates if HA network security is enabled.


|node.name
|link:#node_name[node.name]
a|Node name


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "node.name": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node_name]
```

```
[.api-collapsible-fifth-title]
node.name

Node name


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0a1115a4205f4a9dd252514078e55796]]
= Update the HA network security configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/ha-network`#
```

```
*Introduced In:* 9.18

Updates the HA network security configuration.

== Optional properties

* 'enabled' - Indicates if HA network security for NVLog traffic is
enabled.

== Related ONTAP commands

* 'security ha-network modify'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|enabled
|boolean
|query
|False
a|Specifies if HA network security is enabled.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates if HA network security is enabled.
```

```
|node.name
|link:#node_name[node.name]
a|Node name


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "node.name": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 70385664
| Modifying the HA network security configuration requires an effective HA
version of ONTAP 9.18.1 or later.

| 70385666
| Failed to get node information for node "+++<name>+++\".+++</name>+++

| 70385667
| Failed to enable encryption on node "+++<name>+++\".+++</name>+++

| 70385671
```

```
| Failed to disable encryption on node "+++<name>+++\".+++</name>+++

| 70385673
| HA network security does not support MetroCluster configurations.

| 70385677
| Failed to modify HA network security because the "enabled" field was not
specified.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#node_name]
[.api-collapsible-fifth-title]
```

```
node.name

Node name


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#security_ha_network]
[.api-collapsible-fifth-title]
security_ha_network

Manages the HA network security configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates if HA network security is enabled.


|node.name
|link:#node_name[node.name]
a|Node name


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and update IPsec configuration

:leveloffset: +1


[[IDc01dcb3c8e91f4ac99b5bde95a8701ef]]
= View and update IPsec configuration



== Overview

The following operations are supported:

* GET to retrieve the IPsec status: GET security/ipsec
* Patch to update IPsec status: PATCH security/ipsec





[[ID9d1325ca3e3f69e8482b4e609d573cda]]
= Retrieve an IPsec configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ipsec`#

*Introduced In:* 9.8
```

```
Retrieves IPsec configuration via REST APIs.

== Related ONTAP commands

* 'security ipsec config show'


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|enabled
|boolean
a|Indicates whether or not IPsec is enabled.


|offload_enabled
|boolean
a|Indicates whether or not IPsec hardware offload is enabled.


|replay_window
|integer
a|Replay window size in packets, where 0 indicates that the relay window
is disabled.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "replay_window": 0
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID3790fcc0e9312982922ce681c6a5d047]]
= Update an IPsec configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

```
block]#`/security/ipsec`#

*Introduced In:* 9.8

Updates IPsec configuration via REST APIs.

== Optional properties

* 'enabled' - Enable IPsec.
* 'replay_window' - Replay window size in packets.
* 'offload_enabled' - Enable IPsec hardware offload.

== Related ONTAP commands

* 'security ipsec config modify'


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether or not IPsec is enabled.


|offload_enabled
|boolean
a|Indicates whether or not IPsec hardware offload is enabled.


|replay_window
|integer
a|Replay window size in packets, where 0 indicates that the relay window
is disabled.


|===


.Example request
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "replay_window": 0
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66256898
| Internal error. Failed to enable IPsec.

| 66256899
| Internal error. Failed to disable IPsec.

| 66256913
| Offload setting changes are not allowed when IPsec is enabled.

| 66256915
| Anti-replay protection and IPsec offload cannot be operational at the
same time.

| 66257007
| Failed to update replay window size.

| 66257199
| IPsec is not supported in the current cluster version.

| 66257496
| IPsec must be enabled to support BGP peer groups with MD5. Disable MD5
first in the BGP peer group.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#ipsec]
[.api-collapsible-fifth-title]
ipsec

Manages IPsec configuration via REST APIs.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
```

```
a|Indicates whether or not IPsec is enabled.


|offload_enabled
|boolean
a|Indicates whether or not IPsec hardware offload is enabled.


|replay_window
|integer
a|Replay window size in packets, where 0 indicates that the relay window
is disabled.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage IPsec security certificates

:leveloffset: +1


[[ID76ab44a8964ce3e0c9a1db5b5f237aaf]]
= Manage IPsec security certificates



== Overview

The following APIs can be used to add/remove/retrieve the IPsec CA
```

```
certificates:

* Creation Post: POST security/ipsec/ca-certificates
* Collection Get: GET security/ipsec/ca-certificates
* Instance Get: GET security/ipsec/ca-certificates/{certificate.uuid}
* Instance Delete: DELETE security/ipsec/ca-
certificates/{certificate.uuid}





[[ID608c0334bfb3f3a7392b1e28528c50a8]]
= Retrieve IPsec CA certificates configured for the cluster and all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ipsec/ca-certificates`#

*Introduced In:* 9.10

Retrieves the collection of IPsec CA certificates configured for cluster
and all SVMs.

== Related ONTAP commands

* `security ipsec ca-certificate show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|scope
|string
|query
|False
a|Filter by scope


|certificate.uuid
```

```
|string
|query
|False
a|Filter by certificate.uuid


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
```

```
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#ipsec_ca_certificate[ipsec_ca_certificate]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#certificate]
[.api-collapsible-fifth-title]
```

```
certificate

IPsec CA certificate UUID


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|Certificate UUID


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#ipsec_ca_certificate]
[.api-collapsible-fifth-title]
ipsec_ca_certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|IPsec CA certificate UUID


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID0bfa8adf9a1cc268cac763cb06f0d356]]
= Add a CA certificate to IPsec

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/ipsec/ca-certificates`#

*Introduced In:* 9.10

Add CA certificate to IPsec. The CA certificate should already be
installed on the cluster prior to adding them to IPsec.
The CA certificate can be installed on the cluster using the
/security/certificates endpoint.
The svm.uuid or svm.name should not be supplied for certificates that have
a scope of cluster.

== Related ONTAP commands

* `security ipsec ca-certificate add`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|IPsec CA certificate UUID


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
```

```
|integer
a|Number of records


|records
|array[link:#ipsec_ca_certificate[ipsec_ca_certificate]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "certificate": {
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
```

```
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66257304
| CA certificate is not installed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
```

```
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate

IPsec CA certificate UUID


[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|Certificate UUID


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#ipsec_ca_certificate]
[.api-collapsible-fifth-title]
ipsec_ca_certificate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|IPsec CA certificate UUID


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDdd348080336bdba5d8fafe838b396a99]]
= Delete a CA certificate with the specified UUID from IPsec

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/ipsec/ca-certificates/{certificate.uuid}`#

*Introduced In:* 9.10
```

```
Deletes the IPsec CA certificate with the specified UUID from IPsec.

== Related ONTAP commands

* `security ipsec ca-certificate remove`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|certificate.uuid
|string
|path
|True
a|UUID of the CA certificate to be deleted from IPsec.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66257298
| CA certificate is not installed for IPsec.

| 66257303
| The CA certificate cannot be removed from IPsec because it is not
```

```
installed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID1e6156c5546a0ec45862ecb7598d4760]]
= Retrieve a CA certificate configured for IPpsec

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ipsec/ca-certificates/{certificate.uuid}`#

*Introduced In:* 9.10

Retrieves a specific CA certificate configured for IPsec.

== Related ONTAP commands

* `security ipsec ca-certificate show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|certificate.uuid
|string
|path
|True
a|UUID of the IPsec CA certificate.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|IPsec CA certificate UUID


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
```

```
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#certificate]
[.api-collapsible-fifth-title]
certificate

IPsec CA certificate UUID


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|Certificate UUID


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

```
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage IPsec policies

:leveloffset: +1


[[IDf6fb1ffbae62d3ead32419b4015bcd8d]]
= Manage IPsec policies



== Overview

The following operations are supported:

* Collection Get: GET security/ipsec/policies
* Creation Post: POST security/ipsec/policies
* Instance Get: GET security/ipsec/policies/uuid
* Instance Patch: PATCH security/ipsec/policies/uuid
```

```
* Instance Delete: DELETE security/ipsec/policies/uuid




[[IDb7d4a3ca57151e2f49447a73aa6397c5]]
= Retrieve IPsec policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ipsec/policies`#

*Introduced In:* 9.8

Retrieves the collection of IPsec policies.

== Related ONTAP commands

* `security ipsec policy show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|authentication_method
|string
|query
|False
a|Filter by authentication_method

* Introduced in: 9.10


|uuid
|string
|query
|False
a|Filter by uuid
```

```
|local_identity
|string
|query
|False
a|Filter by local_identity


|protocol
|string
|query
|False
a|Filter by protocol


|name
|string
|query
|False
a|Filter by name

* maxLength: 64
* minLength: 1


|certificate_modify_keepsa
|boolean
|query
|False
a|Filter by certificate_modify_keepsa

* Introduced in: 9.18


|remote_endpoint.netmask
|string
|query
|False
a|Filter by remote_endpoint.netmask


|remote_endpoint.port
|string
|query
|False
a|Filter by remote_endpoint.port
```

```
|remote_endpoint.address
|string
|query
|False
a|Filter by remote_endpoint.address


|remote_endpoint.family
|string
|query
|False
a|Filter by remote_endpoint.family


|remote_identity
|string
|query
|False
a|Filter by remote_identity


|enabled
|boolean
|query
|False
a|Filter by enabled


|certificate.name
|string
|query
|False
a|Filter by certificate.name

* Introduced in: 9.10


|certificate.uuid
|string
|query
|False
a|Filter by certificate.uuid

* Introduced in: 9.10
```

```
|local_endpoint.netmask
|string
|query
|False
a|Filter by local_endpoint.netmask


|local_endpoint.port
|string
|query
|False
a|Filter by local_endpoint.port


|local_endpoint.address
|string
|query
|False
a|Filter by local_endpoint.address


|local_endpoint.family
|string
|query
|False
a|Filter by local_endpoint.family


|scope
|string
|query
|False
a|Filter by scope


|ipspace.name
|string
|query
|False
a|Filter by ipspace.name


|ipspace.uuid
|string
|query
|False
a|Filter by ipspace.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|ppk.identity
|string
|query
|False
a|Filter by ppk.identity

* Introduced in: 9.17
* maxLength: 64
* minLength: 6


|action
|string
|query
|False
a|Filter by action

* Introduced in: 9.15


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
```

```
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
```

```
      "action": "string",
      "authentication_method": "string",
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "local_endpoint": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24",
        "port": "23"
      },
      "local_identity": "string",
      "name": "string",
      "ppk": {
        "identity": "string"
      },
      "protocol": "17",
      "remote_endpoint": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24",
        "port": "23"
      },
      "remote_identity": "string",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
```

```
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#certificate]
[.api-collapsible-fifth-title]
certificate

Certificate for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===
```

```
[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#local_endpoint]
[.api-collapsible-fifth-title]
local_endpoint

Local endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
```

```
a|IPv4 or IPv6 address


|family
|string
a|IPv4 or IPv6


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===


[#ppk]
[.api-collapsible-fifth-title]
ppk

Post-quantum pre-shared key information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|identity
|string
a|Post-quantum pre-shared key identity.


|===


[#remote_endpoint]
[.api-collapsible-fifth-title]
remote_endpoint
```

```
Remote endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|family
|string
a|IPv4 or IPv6


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#records]
[.api-collapsible-fifth-title]
records

IPsec policy object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec policy.


|authentication_method
|string
a|Authentication method for the IPsec policy.


|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.
```

```
|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|name
|string
a|IPsec policy name.


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.
```

```
|remote_identity
|string
a|Remote Identity


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|Unique identifier of the IPsec policy.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDce8425f35be544327813f198b09731fd]]
= Create an IPsec policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/ipsec/policies`#

*Introduced In:* 9.8

Creates an IPsec policy.

== Related ONTAP commands

* `security ipsec policy create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.
```

```
* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec policy.


|authentication_method
|string
a|Authentication method for the IPsec policy.


|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


|local_endpoint
```

```
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|name
|string
a|IPsec policy name.


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|secret_key
|string
a|Pre-shared key for IKE negotiation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|Unique identifier of the IPsec policy.
```

```
|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "action": "string",
  "authentication_method": "string",
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "ipspace": {
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24",
    "port": "23"
  },
  "local_identity": "string",
  "name": "string",
  "ppk": {
    "identity": "string",
    "shared_key": "string"
  },
  "protocol": "17",
  "remote_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24",
    "port": "23"
  },
  "remote_identity": "string",
  "secret_key": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====
```

```
== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "action": "string",
      "authentication_method": "string",
      "certificate": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "ipspace": {
        "name": "Default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "local_endpoint": {
        "address": "10.10.10.7",
        "netmask": "24",
        "port": "23"
      },
```

```
      "local_identity": "string",
      "name": "string",
      "ppk": {
        "identity": "string",
        "shared_key": "string"
      },
      "protocol": "17",
      "remote_endpoint": {
        "address": "10.10.10.7",
        "netmask": "24",
        "port": "23"
      },
      "remote_identity": "string",
      "secret_key": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66257099
| Only one protocol can be specified.

| 66257100
| Only one local port can be specified.

| 66257101
| Only one remote port can be specified.

| 66257102
| Internet Key Exchange version 1 (IKEv1) is not supported.

| 66257104
| IPsec policy with same name already exists in this SVM.

| 66257107
| The specified pre-shared key is not a valid hexadecimal string.

| 66257109
| The specified pre-shared key is not a valid Base64 encoded binary
string.

| 66257110
| Failed to a create policy sequencing value.

| 66257112
| The IPsec policy with the action specified does not provide packet
protection and the authentication method provided for the policy will be
ignored.

| 66257113
| Only one local IP subnet can be specified.

| 66257114
| Only one remote IP subnet can be specified.

| 66257115
| Port ranges containing more than one port are not supported.

| 66257117
| IPsec is not supported on the SVM specified in the policy, IPsec is
supported on data SVMs only.
```

```
| 66257120
| The subnet selector must be a host address (An IPv4 address with a 32-
bit netmask or an IPv6 address with a 128-bit netmask).

| 66257121
| The maximum limit of IPsec policies has been reached for the specified
SVM.

| 66257125
| The local_endpoint.address must be specified with
local_endpoint.netmask.

| 66257126
| The remote_endpoint.address must be specified with
remote_endpoint.netmask.

| 66257127
| The local subnet must be configured as a non-zero subnet.

| 66257128
| Invalid ANY wildcard subnet.

| 66257129
| A specific local or remote port number is required when the remote
subnet is configured as an ANY wildcard subnet.

| 66257130
| The maximum limit of IPsec policies has been reached for the cluster.

| 66257131
| ESP in UDPv6 Encapsulation is not supported.

| 66257132
| Invalid value for port field. Value should be in range <0-65535>.

| 66257133
| A pre-shared key is needed for the PSK authentication method. Use the
secret_key option to specify a key.

| 66257134
| An end-entity certificate is needed for the PKI authentication method.
Use the certificate.uuid option to specify an end-entity certificate.

| 66257135
| The specified certificate is not found on the SVM.
```

```
| 66257136
| A certificate is not needed for the PSK authentication method.

| 66257137
| A pre-shared key is not needed for the PKI authentication method.

| 66257138
| Remote identity is required when using certificate verification.

| 66257139
| Certificate with the specified UUID was not found.

| 66257140
| Only certificates with a client or server type are supported.

| 66257142
| Failed to create IPsec policy because the specified SVM is being
migrated.

| 66257143
| Invalid IPsec policy provided. The subnet must be non-empty.

| 66257144
| The IPsec policy actions ESP TRANSPORT and ESP UDP each provide packet
protection and requires a secret key or certificate for authentication.

| 66257148
| The policy name does not meet required ASCII-range characters length.

| 66257153
| The policy requires a post-quantum pre-shared key because ppk-identity
is specified for this policy.

| 66257154
| The policy requires a post-quantum pre-shared identity because ppk-
shared-key is specified for this policy.

| 66257199
| Not all of the nodes in the cluster are running a version that supports
the IPsec feature.

| 66257200
| The shared key does not meet required ASCII-range characters length.

| 66257201
| Support for the feature available with effective cluster version or
later.
```

```
| 66257202
| The specified SVM name is invalid.

| 66257203
| The specified SVM UUID is invalid.

| 66257204
| The specified IPspace UUID and IPspace name refer to different IPspaces.

| 66257205
| The specified SVM must exist in the specified IPspace.

| 66257305
| The certificate UUID does not match the provided certificate name.

| 66257396
| IPsec is not supported for the admin SVM in a MetroCluster
configuration.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate

Certificate for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#local_endpoint]
[.api-collapsible-fifth-title]
local_endpoint

Local endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===


[#ppk]
[.api-collapsible-fifth-title]
```

```
ppk

Post-quantum pre-shared key information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|identity
|string
a|Post-quantum pre-shared key identity.


|shared_key
|string
a|Post-quantum pre-shared key.


|===


[#remote_endpoint]
[.api-collapsible-fifth-title]
remote_endpoint

Remote endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.
```

```
|port
|string
a|Application port to be covered by the IPsec policy


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#ipsec_policy]
[.api-collapsible-fifth-title]
ipsec_policy

IPsec policy object.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|action
|string
a|Action for the IPsec policy.


|authentication_method
|string
a|Authentication method for the IPsec policy.


|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|name
|string
```

```
a|IPsec policy name.


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|secret_key
|string
a|Pre-shared key for IKE negotiation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|Unique identifier of the IPsec policy.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#records]
[.api-collapsible-fifth-title]
```

```
records

IPsec policy object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec policy.


|authentication_method
|string
a|Authentication method for the IPsec policy.


|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.
```

```
|local_identity
|string
a|Local Identity


|name
|string
a|IPsec policy name.


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|secret_key
|string
a|Pre-shared key for IKE negotiation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|Unique identifier of the IPsec policy.


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDcc510e78d9a094a37393037c97d3f0c5]]
= Delete an IPsec policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/ipsec/policies/{uuid}`#

*Introduced In:* 9.8
```

```
Deletes a specific IPsec policy.

== Related ONTAP commands

* `security ipsec policy delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|IPsec policy UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66257096
| Internal error. Failed to purge connections associated with the IPsec
policy.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb6825db2836145cae4fe65e228b3038b]]
= Retrieve an IPsec policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ipsec/policies/{uuid}`#

*Introduced In:* 9.8

Retrieves a specific IPsec policy.

== Related ONTAP commands

* `security ipsec policy show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|IPsec policy UUID


|fields
|array[string]
```

```
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec policy.


|authentication_method
|string
a|Authentication method for the IPsec policy.


|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.
```

```
|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|name
|string
a|IPsec policy name.


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.
```

```
|uuid
|string
a|Unique identifier of the IPsec policy.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "action": "string",
  "authentication_method": "string",
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_endpoint": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24",
    "port": "23"
  },
  "local_identity": "string",
  "name": "string",
  "ppk": {
    "identity": "string"
  },
  "protocol": "17",
```

```
  "remote_endpoint": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24",
    "port": "23"
  },
  "remote_identity": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#certificate]
[.api-collapsible-fifth-title]
certificate

Certificate for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#local_endpoint]
[.api-collapsible-fifth-title]
local_endpoint

Local endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|family
|string
a|IPv4 or IPv6


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
```

```
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===


[#ppk]
[.api-collapsible-fifth-title]
ppk

Post-quantum pre-shared key information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|identity
|string
a|Post-quantum pre-shared key identity.


|===


[#remote_endpoint]
[.api-collapsible-fifth-title]
remote_endpoint

Remote endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
```

```
|string
a|IPv4 or IPv6 address


|family
|string
a|IPv4 or IPv6


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
```

```
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDac2b4ab4aebd6d40b6043d7ae77533d5]]
= Update an IPsec policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/ipsec/policies/{uuid}`#

*Introduced In:* 9.8

Updates a specific IPsec policy.

== Related ONTAP commands

* `security ipsec policy modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|IPsec policy UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|ppk
|link:#ppk[ppk]
```

```
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|uuid
|string
a|Unique identifier of the IPsec policy.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "certificate": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24",
    "port": "23"
  },
  "local_identity": "string",
  "ppk": {
    "identity": "string"
  },
  "protocol": "17",
  "remote_endpoint": {
    "address": "10.10.10.7",
```

```
    "netmask": "24",
    "port": "23"
  },
  "remote_identity": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66257097
| Internal error. Failed to update the IPsec policy.

| 66257099
| Only one protocol can be specified.

| 66257100
| Only one local port can be specified.

| 66257101
| Only one remote port can be specified.

| 66257110
| Failed to create a policy sequencing value.

| 66257113
| Only one local IP subnet can be specified.

| 66257114
| Only one remote IP subnet can be specified.

| 66257115
| Port ranges containing more than one port are not supported.
```

```
| 66257120
| The subnet selector must be a host address (An IPv4 address with a 32-
bit netmask or an IPv6 address with a 128-bit netmask).

| 66257125
| The local_endpoint.address must be specified with
local_endpoint.netmask.

| 66257126
| The remote_endpoint.address must be specified with
remote_endpoint.netmask.

| 66257127
| The local subnet must be configured as a non-zero subnet.

| 66257128
| Invalid ANY wildcard subnet.

| 66257129
| A specific local or remote port number is required when the remote
subnet is configured as an ANY wildcard subnet.

| 66257131
| ESP in UDPv6 Encapsulation is not supported.

| 66257135
| The specified certificate is not found on the SVM.

| 66257136
| A certificate is not needed for the PSK authentication method.

| 66257138
| Remote identity is required when using certificate verification.

| 66257139
| Certificate with the specified UUID was not found.

| 66257140
| Only certificates with a client or server type are supported.

| 66257142
| Failed to create IPsec policy because the specified SVM is being
migrated.

| 66257143
| Invalid IPsec policy provided. The subnet must be non-empty.
```

```
| 66257153
| The policy requires a post-quantum pre-shared key because ppk-identity
is specified for this policy.

| 66257154
| The policy requires a post-quantum pre-shared identity because ppk-
shared-key is specified for this policy.

| 66257199
| Not all of the nodes in the cluster are running a version that supports
the IPsec feature.

| 66257200
| The shared key does not meet required ASCII-range characters length.

| 66257201
| Support for the feature available with effective cluster version or
later.

| 66257305
| The certificate UUID does not match the provided certificate name.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#certificate]
[.api-collapsible-fifth-title]
certificate

Certificate for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#local_endpoint]
[.api-collapsible-fifth-title]
local_endpoint

Local endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
|string
a|Application port to be covered by the IPsec policy


|===
```

```
[#ppk]
[.api-collapsible-fifth-title]
ppk

Post-quantum pre-shared key information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|identity
|string
a|Post-quantum pre-shared key identity.


|===


[#remote_endpoint]
[.api-collapsible-fifth-title]
remote_endpoint

Remote endpoint for the IPsec policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, the
default value is 64 with a valid range of 1 to 127. Output is always the
netmask length.


|port
```

```
|string
a|Application port to be covered by the IPsec policy


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#ipsec_policy]
[.api-collapsible-fifth-title]
ipsec_policy

IPsec policy object.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|certificate
|link:#certificate[certificate]
a|Certificate for the IPsec policy.


|certificate_modify_keepsa
|boolean
a|Indicates whether old security associations are kept upon certificate
modification.


|enabled
|boolean
a|Indicates whether or not the policy is enabled.


|local_endpoint
|link:#local_endpoint[local_endpoint]
a|Local endpoint for the IPsec policy.


|local_identity
|string
a|Local Identity


|ppk
|link:#ppk[ppk]
a|Post-quantum pre-shared key information.


|protocol
|string
a|Lower layer protocol to be covered by the IPsec policy.


|remote_endpoint
|link:#remote_endpoint[remote_endpoint]
a|Remote endpoint for the IPsec policy.


|remote_identity
|string
a|Remote Identity


|uuid
```

```
|string
a|Unique identifier of the IPsec policy.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View IPsec security associations

:leveloffset: +1


[[IDc556653e80fa0875fdb57e1b38015a25]]
= View IPsec security associations



== Overview

* Collection Get: GET security/ipsec/security-associations
* Instance Get: GET security/ipsec/security-associations/uuid





[[ID27b055031aa29f254b9b01b977086656]]
= Retrieve IPsec and IKE security associations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/security/ipsec/security-associations`#

*Introduced In:* 9.8

Retrieves the IPsec and IKE (Internet Key Exchange) security associations.

== Related ONTAP commands

* `security ipsec show-ipsecsa`
* `security ipsec show-ikesa`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|type
|string
|query
|False
a|Filter by type


|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|ike.is_initiator
|boolean
|query
```

```
|False
a|Filter by ike.is_initiator


|ike.initiator_security_parameter_index
|string
|query
|False
a|Filter by ike.initiator_security_parameter_index


|ike.state
|string
|query
|False
a|Filter by ike.state


|ike.responder_security_parameter_index
|string
|query
|False
a|Filter by ike.responder_security_parameter_index


|ike.version
|integer
|query
|False
a|Filter by ike.version


|ike.authentication
|string
|query
|False
a|Filter by ike.authentication


|remote_address
|string
|query
|False
a|Filter by remote_address


|cipher_suite
```

```
|string
|query
|False
a|Filter by cipher_suite


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|uuid
|string
|query
|False
a|Filter by uuid


|ipsec.action
|string
|query
|False
a|Filter by ipsec.action


|ipsec.state
|string
|query
|False
a|Filter by ipsec.state


|ipsec.inbound.offload_bytes
|integer
|query
|False
a|Filter by ipsec.inbound.offload_bytes
```

```
* Introduced in: 9.16


|ipsec.inbound.packets
|integer
|query
|False
a|Filter by ipsec.inbound.packets


|ipsec.inbound.security_parameter_index
|string
|query
|False
a|Filter by ipsec.inbound.security_parameter_index


|ipsec.inbound.bytes
|integer
|query
|False
a|Filter by ipsec.inbound.bytes


|ipsec.inbound.offload_packets
|integer
|query
|False
a|Filter by ipsec.inbound.offload_packets

* Introduced in: 9.16


|ipsec.outbound.security_parameter_index
|string
|query
|False
a|Filter by ipsec.outbound.security_parameter_index


|ipsec.outbound.packets
|integer
|query
|False
a|Filter by ipsec.outbound.packets
```

```
|ipsec.outbound.offload_bytes
|integer
|query
|False
a|Filter by ipsec.outbound.offload_bytes

* Introduced in: 9.16


|ipsec.outbound.offload_packets
|integer
|query
|False
a|Filter by ipsec.outbound.offload_packets

* Introduced in: 9.16


|ipsec.outbound.bytes
|integer
|query
|False
a|Filter by ipsec.outbound.bytes


|scope
|string
|query
|False
a|Filter by scope


|local_address
|string
|query
|False
a|Filter by local_address


|lifetime
|integer
|query
|False
a|Filter by lifetime


|policy_name
```

```
|string
|query
|False
a|Filter by policy_name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
```

```
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
      "cipher_suite": "string",
      "ike": {
        "authentication": "string",
        "initiator_security_parameter_index": "string",
        "responder_security_parameter_index": "string",
        "state": "string"
      },
      "ipsec": {
        "action": "string",
        "inbound": {
          "security_parameter_index": "string"
        },
        "outbound": {
          "security_parameter_index": "string"
        },
        "state": "string"
      },
      "local_address": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "policy_name": "string",
      "remote_address": "string",
      "scope": "string",
```

```
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
```

```
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
```

```
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|===


[#ike]
[.api-collapsible-fifth-title]
ike

Objects containing parameters specific to IKE (Internet Key Exchange)
security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication
|string
a|Authentication method for internet key exchange protocol.


|initiator_security_parameter_index
|string
a|Initiator's security parameter index for the IKE security association.


|is_initiator
|boolean
a|Indicates whether or not IKE has been initiated by this node.


|responder_security_parameter_index
|string
a|Responder's security parameter index for the IKE security association.


|state
|string
a|State of the IKE connection.
```

```
|version
|integer
a|Internet key exchange protocol version.


|===


[#inbound]
[.api-collapsible-fifth-title]
inbound

Status for inbound parameters for the IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes
|integer
a|Number of inbound bytes for the IPsec security association.


|offload_bytes
|integer
a|Number of inbound bytes processed by offload for the IPsec security
association.


|offload_packets
|integer
a|Number of inbound packets processed by offload for the IPsec security
association.


|packets
|integer
a|Number of inbound packets for the IPsec security association.


|security_parameter_index
|string
a|Inbound security parameter index for the IPSec security association.
```

```
|===


[#outbound]
[.api-collapsible-fifth-title]
outbound

Status for outbound parameters for the IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes
|integer
a|Number of outbound bytes for the IPsec security association.


|offload_bytes
|integer
a|Number of outbound bytes processed by offload for the IPsec security
association.


|offload_packets
|integer
a|Number of outbound packets processed by offload for the IPsec security
association.


|packets
|integer
a|Number of outbound packets for the IPsec security association.


|security_parameter_index
|string
a|Outbound security parameter index for the IPSec security association.


|===
```

```
[#ipsec]
[.api-collapsible-fifth-title]
ipsec

Objects containing parameters specific to IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec security association.


|inbound
|link:#inbound[inbound]
a|Status for inbound parameters for the IPsec security association.


|outbound
|link:#outbound[outbound]
a|Status for outbound parameters for the IPsec security association.


|state
|string
a|State of the IPsec security association.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

Node with the security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#records]
[.api-collapsible-fifth-title]
records

Security association object for IPsec security association and IKE
(Internet Key Exchange) security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cipher_suite
|string
a|Cipher suite for the security association.


|ike
|link:#ike[ike]
a|Objects containing parameters specific to IKE (Internet Key Exchange)
security association.


|ipsec
|link:#ipsec[ipsec]
a|Objects containing parameters specific to IPsec security association.
```

```
|lifetime
|integer
a|Lifetime for the security association in seconds.


|local_address
|string
a|Local address of the security association.


|node
|link:#node[node]
a|Node with the security association.


|policy_name
|string
a|Policy name for the security association.


|remote_address
|string
a|Remote address of the security association.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of security association, it can be IPsec or IKE (Internet Key
Exchange).


|uuid
|string
a|Unique identifier of the security association.
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID210129777b7e76c37852d636f11e5686]]
= Retrieve an IPsec or IKE security association

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/security/ipsec/security-associations/{uuid}`#

*Introduced In:* 9.8

Retrieves a specific IPsec or IKE (Internet Key Exchange) security
association.

== Related ONTAP commands

* `security ipsec show-ipsecsa`
* `security ipsec show-ikesa`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of IPsec or IKE security association.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|cipher_suite
|string
a|Cipher suite for the security association.


|ike
|link:#ike[ike]
a|Objects containing parameters specific to IKE (Internet Key Exchange)
security association.


|ipsec
|link:#ipsec[ipsec]
a|Objects containing parameters specific to IPsec security association.


|lifetime
|integer
a|Lifetime for the security association in seconds.


|local_address
|string
a|Local address of the security association.


|node
|link:#node[node]
a|Node with the security association.


|policy_name
|string
a|Policy name for the security association.


|remote_address
|string
a|Remote address of the security association.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
```

```
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of security association, it can be IPsec or IKE (Internet Key
Exchange).


|uuid
|string
a|Unique identifier of the security association.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "cipher_suite": "string",
  "ike": {
    "authentication": "string",
    "initiator_security_parameter_index": "string",
    "responder_security_parameter_index": "string",
    "state": "string"
  },
  "ipsec": {
    "action": "string",
    "inbound": {
      "security_parameter_index": "string"
    },
    "outbound": {
      "security_parameter_index": "string"
    },
    "state": "string"
  },
  "local_address": "string",
  "node": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "policy_name": "string",
  "remote_address": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#ike]
[.api-collapsible-fifth-title]
ike

Objects containing parameters specific to IKE (Internet Key Exchange)
security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication
|string
a|Authentication method for internet key exchange protocol.


|initiator_security_parameter_index
|string
a|Initiator's security parameter index for the IKE security association.


|is_initiator
```

```
|boolean
a|Indicates whether or not IKE has been initiated by this node.


|responder_security_parameter_index
|string
a|Responder's security parameter index for the IKE security association.


|state
|string
a|State of the IKE connection.


|version
|integer
a|Internet key exchange protocol version.


|===


[#inbound]
[.api-collapsible-fifth-title]
inbound

Status for inbound parameters for the IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes
|integer
a|Number of inbound bytes for the IPsec security association.


|offload_bytes
|integer
a|Number of inbound bytes processed by offload for the IPsec security
association.


|offload_packets
```

```
|integer
a|Number of inbound packets processed by offload for the IPsec security
association.


|packets
|integer
a|Number of inbound packets for the IPsec security association.


|security_parameter_index
|string
a|Inbound security parameter index for the IPSec security association.


|===


[#outbound]
[.api-collapsible-fifth-title]
outbound

Status for outbound parameters for the IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes
|integer
a|Number of outbound bytes for the IPsec security association.


|offload_bytes
|integer
a|Number of outbound bytes processed by offload for the IPsec security
association.


|offload_packets
|integer
a|Number of outbound packets processed by offload for the IPsec security
association.
```

```
|packets
|integer
a|Number of outbound packets for the IPsec security association.


|security_parameter_index
|string
a|Outbound security parameter index for the IPSec security association.


|===


[#ipsec]
[.api-collapsible-fifth-title]
ipsec

Objects containing parameters specific to IPsec security association.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|string
a|Action for the IPsec security association.


|inbound
|link:#inbound[inbound]
a|Status for inbound parameters for the IPsec security association.


|outbound
|link:#outbound[outbound]
a|Status for outbound parameters for the IPsec security association.


|state
|string
a|State of the IPsec security association.
```

```
|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

Node with the security association.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage JIT privilege user configurations

:leveloffset: +1


[[IDe4b80951c141e81b81a16ee951e216a9]]
= Manage JIT privilege user configurations



== Overview

This API is used to create, retrieve, and delete relevant information
related to user-specific JIT privilege session configurations.
Prerequisites:
  You must create a user that has a JIT supported login method configured
on the SVM that you want to create the elevated session configurations.

== Examples

=== Creating the JIT session configuration for a user

----

# The API:
POST "/api/security/jit-privilege-users"

# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/security/jit-privilege-
users?return_records=true" -H "accept: application/hal+json" -d '{
"owner": { "name": "Cserver"} , "account": { "name": "joe"},
"application": "ssh","role": {"name": "admin"},"session_validity":
"PT30M","jit_validity": "PT2H10M20S"}'

# The response:
{}
----

=== Retrieving the configured JIT privilege sessions for users

Retrieves the JIT privilege user session configurations or a filtered list
(for a specific SVM, user, etc).

----

# The API:
GET "/api/security/jit-privilege-users"

# The call:
curl -X GET  "https://<mgmt-ip>/api/security/jit-privilege-users" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "owner": {
        "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
        "name": "Cserver",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
          }
        }
      },
      "account": {
        "name": "testUser"
      },
      "application": "ssh",
        "_links": {
          "self": {
            "href": "/api/security/jit-privilege-users/b009a9e7-4081-b576-
7575-ada21efcaf16/testUser/ssh"
          }
```

```
        }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/jit-privilege-users"
    }
  }
}
```


```
[[ID60de63b296a03d4087d5f86cff1ad808]]
= Retrieve the JIT privilege user configurations for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/jit-privilege-users`#

*Introduced In:* 9.17

Retrieves the JIT privilege user configurations for an SVM.

== Related ONTAP commands

* `security jit-privilege user show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|start_time
|string
|query
|False
a|Filter by start_time
```

```
|role.name
|string
|query
|False
a|Filter by role.name


|jit_state
|string
|query
|False
a|Filter by jit_state


|account.name
|string
|query
|False
a|Filter by account.name


|application
|string
|query
|False
a|Filter by application


|jit_validity
|string
|query
|False
a|Filter by jit_validity


|session_validity
|string
|query
|False
a|Filter by session_validity


|end_time
|string
|query
```

```
|False
a|Filter by end_time


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|comment
|string
|query
|False
a|Filter by comment


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
```

```
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_jit_privilege_user[security_jit_privilege_user]]
```

```
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "account": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "joe.smith"
      },
      "application": "string",
      "comment": "Comment text",
      "end_time": "2024-11-08 11:15:31 +0000",
      "jit_state": "active",
      "jit_validity": "P90D",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "role": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "session_validity": "PT1H",
      "start_time": "2024-11-08 10:15:31 +0000"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
```

```
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

The name of the user.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner
```

```
Owner name and UUID that uniquely identifies the JIT settings for a user.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role]
[.api-collapsible-fifth-title]
role

The role the user will be elevated to.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Role name
```

```
|===


[#security_jit_privilege_user]
[.api-collapsible-fifth-title]
security_jit_privilege_user

[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|The name of the user.


|application
|string
a|The name of the application.


|comment
|string
a|The description of the elevation rule.


|end_time
|string
a|The end date and time of the JIT privilege session in UTC.


|jit_state
|string
a|The current status of the rule: preactive, idle, or active.


|jit_validity
|string
a|The JIT validity period on this SVM for this user.


|owner
|link:#owner[owner]
```

```
a|Owner name and UUID that uniquely identifies the JIT settings for a
user.


|role
|link:#role[role]
a|The role the user will be elevated to.


|session_validity
|string
a|The session validity period on this SVM for this user.


|start_time
|string
a|The start date and time of the JIT privilege session in UTC.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf67904ab6d0cd82a8ac6f5960ef493d2]]
= Create JIT privilege user configurations for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/jit-privilege-users`#

*Introduced In:* 9.17

Creates the JIT privilege user configurations for an SVM.

== Required properties
```

```
* `owner.uuid` - Account owner UUID.
* `account.name` - User name.
* `application` - Application

== Related ONTAP commands

* `security jit-privilege user create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|The name of the user.


|application
```

```
|string
a|The name of the application.


|comment
|string
a|The description of the elevation rule.


|end_time
|string
a|The end date and time of the JIT privilege session in UTC.


|jit_state
|string
a|The current status of the rule: preactive, idle, or active.


|jit_validity
|string
a|The JIT validity period on this SVM for this user.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the JIT settings for a
user.


|role
|link:#role[role]
a|The role the user will be elevated to.


|session_validity
|string
a|The session validity period on this SVM for this user.


|start_time
|string
a|The start date and time of the JIT privilege session in UTC.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "account": {
    "name": "joe.smith"
  },
  "application": "string",
  "comment": "Comment text",
  "end_time": "2024-11-08 11:15:31 +0000",
  "jit_state": "active",
  "jit_validity": "P90D",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "role": {
    "name": "admin"
  },
  "session_validity": "PT1H",
  "start_time": "2024-11-08 10:15:31 +0000"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#security_jit_privilege_user[security_jit_privilege_user]]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "account": {
        "name": "joe.smith"
      },
      "application": "string",
      "comment": "Comment text",
      "end_time": "2024-11-08 11:15:31 +0000",
      "jit_state": "active",
      "jit_validity": "P90D",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "role": {
        "name": "admin"
      },
      "session_validity": "PT1H",
      "start_time": "2024-11-08 10:15:31 +0000"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
```

```
|string
//end row
//end table
|===
```

```
== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#account_reference]
[.api-collapsible-fifth-title]
account_reference

The name of the user.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the JIT settings for a user.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role]
[.api-collapsible-fifth-title]
role

The role the user will be elevated to.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|Role name


|===


[#security_jit_privilege_user]
[.api-collapsible-fifth-title]
security_jit_privilege_user

[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|The name of the user.


|application
|string
a|The name of the application.


|comment
|string
a|The description of the elevation rule.


|end_time
|string
a|The end date and time of the JIT privilege session in UTC.


|jit_state
|string
a|The current status of the rule: preactive, idle, or active.
```

```
|jit_validity
|string
a|The JIT validity period on this SVM for this user.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the JIT settings for a
user.


|role
|link:#role[role]
a|The role the user will be elevated to.


|session_validity
|string
a|The session validity period on this SVM for this user.


|start_time
|string
a|The start date and time of the JIT privilege session in UTC.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


[[IDd65a81212dd4f232c562285db8ab7180]]
```

```
= Delete JIT privilege user configurations for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/jit-privilege-
users/{owner.uuid}/{account.name}/{application}`#

*Introduced In:* 9.17

Deletes the JIT privilege user configurations for an SVM.

== Related ONTAP commands

* `security jit-privilege user delete`

== Required properties

* `owner.uuid`
* `application`
* `account.name`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|account.name
|string
|path
|True
a|

|application
|string
|path
```

```
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDbe7b4c9e79dbcb4dfe2c87a3abbf99a1]]
= Retrieve the JIT privilege user configurations for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/jit-privilege-
users/{owner.uuid}/{account.name}/{application}`#

*Introduced In:* 9.17

Retrieves the JIT privilege user configurations for an SVM.

== Related ONTAP commands

* `security jit-privilege user show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|owner.uuid
|string
|path
|True
a|Used to identify a cluster or an SVM


|account.name
|string
|path
|True
a|The username.


|application
|string
|path
|True
a|The name of the application.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|The name of the user.


|application
|string
```

```
a|The name of the application.


|comment
|string
a|The description of the elevation rule.


|end_time
|string
a|The end date and time of the JIT privilege session in UTC.


|jit_state
|string
a|The current status of the rule: preactive, idle, or active.


|jit_validity
|string
a|The JIT validity period on this SVM for this user.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the JIT settings for a
user.


|role
|link:#role[role]
a|The role the user will be elevated to.


|session_validity
|string
a|The session validity period on this SVM for this user.


|start_time
|string
a|The start date and time of the JIT privilege session in UTC.


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "account": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "joe.smith"
  },
  "application": "string",
  "comment": "Comment text",
  "end_time": "2024-11-08 11:15:31 +0000",
  "jit_state": "active",
  "jit_validity": "P90D",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "session_validity": "PT1H",
  "start_time": "2024-11-08 10:15:31 +0000"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

The name of the user.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account
```

```
|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the JIT settings for a user.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role]
[.api-collapsible-fifth-title]
role

The role the user will be elevated to.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|Role name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



= Manage global JIT privilege configurations

:leveloffset: +1


[[IDa339295072f256d07eb0e285549b7a21]]
= Manage global JIT privilege configurations



== Overview

This API is used to retrieve and configure relevant information related to
the Global JIT (Just In Time) privilege configuration in the cluster.

== Examples

=== Retrieving the configured JIT privilege settings for the cluster

Retrieves the JIT privileges for the cluster or a filtered list (for a
```

```
specific SVM).

----

# The API:
GET "/api/security/jit-privileges"

# The call
curl -k https://<mgmt-ip>/api/security/jit-privileges

# the response:
{
"records": [
{
  "owner": {
    "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "Cserver",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
  "application": "ssh",
  "_links": {
    "self": {
      "href": "/api/security/jit-privileges/b009a9e7-4081-b576-7575-
ada21efcaf16/ssh"
    }
  }
}
]
}
----




[[ID2d7143da2d7222249d1fbd7db31eb032]]
= Retrieve global JIT privilege configurations for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/jit-privileges`#

*Introduced In:* 9.17
```

```
Retrieves global JIT privilege configurations on an SVM.

== Related ONTAP commands

* `security jit-privilege show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|default_session_validity_period
|string
|query
|False
a|Filter by default_session_validity_period


|max_jit_validity_period
|string
|query
|False
a|Filter by max_jit_validity_period


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|application
```

```
|string
|query
|False
a|Filter by application


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
```

```
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_jit_privilege[security_jit_privilege]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
```

```
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": "string",
      "default_session_validity_period": "PT1H",
      "max_jit_validity_period": "P90D",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the JIT privilege global
settings.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_jit_privilege]
[.api-collapsible-fifth-title]
security_jit_privilege

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application
|string
a|The name of the application.


|default_session_validity_period
|string
a|The default session validity period on this SVM.


|max_jit_validity_period
|string
a|The maximum JIT validity period allowed on this SVM.


|owner
|link:#owner[owner]
```

```
a|Owner name and UUID that uniquely identifies the JIT privilege global
settings.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


[[IDad63936fab345203b23d6fe584502de4]]
= Retrieve JIT privilege configurations for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/jit-privileges/{owner.uuid}/{application}`#

*Introduced In:* 9.17

Retrieves the JIT privilege configurations for an SVM.

== Related ONTAP commands

* `security jit-privilege show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|owner.uuid
|string
|path
|True
a|Used to identify a cluster or an SVM.


|application
|string
|path
|True
a|Application name.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|application
|string
a|The name of the application.


|default_session_validity_period
|string
```

```
a|The default session validity period on this SVM.


|max_jit_validity_period
|string
a|The maximum JIT validity period allowed on this SVM.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the JIT privilege global
settings.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "application": "string",
  "default_session_validity_period": "PT1H",
  "max_jit_validity_period": "P90D",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the JIT privilege global
settings.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

```
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDfbd458aa45e5dcef91663b088eff695d]]
= Modify JIT privilege configurations for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/jit-privileges/{owner.uuid}/{application}`#

*Introduced In:* 9.17

Modifies the JIT privilege configurations for an SVM.

== Related ONTAP commands

* `security jit-privilege modify`


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Used to identify a cluster or an SVM.


|application
|string
|path
|True
a|Application name.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_session_validity_period
|string
a|The default session validity period on this SVM.


|max_jit_validity_period
|string
a|The maximum JIT validity period allowed on this SVM.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "default_session_validity_period": "PT1H",
  "max_jit_validity_period": "P90D"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_session_validity_period
|string
a|The default session validity period on this SVM.


|max_jit_validity_period
|string
a|The maximum JIT validity period allowed on this SVM.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "default_session_validity_period": "PT1H",
  "max_jit_validity_period": "P90D"
}
====

== Error
```

Status: Default, Error

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the JIT privilege global
settings.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
```

```
PATCH method.


|===


[#security_jit_privilege]
[.api-collapsible-fifth-title]
security_jit_privilege

[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_session_validity_period
|string
a|The default session validity period on this SVM.


|max_jit_validity_period
|string
a|The maximum JIT validity period allowed on this SVM.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



= View and update key manager configurations

:leveloffset: +1
```

```
[[ID2f056c51a7267c2f079ebd1830263061]]
= View and update key manager configurations



== Overview

Retrieves or modifies the key management configuration options. The
following operations are supported:

* GET
* PATCH

== Examples

=== Retrieving cluster-level key manager configurations

The following example shows how to retrieve cluster-level manager
configurations.

----

# The API:
GET /api/security/key-manager-configs

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-manager-configs' -H
'accept: application/hal+json'

# The response:
{
"cc_mode_enabled": false,
"health_monitor_polling_interval": 15,
"cloud_kms_retry_count": 3,
"health_monitor_policy": {
  "akv": {
    "enabled": true,
    "manage_volume_offline": true
  },
  "gcp": {
    "enabled": true,
    "manage_volume_offline": true
  },
  "aws": {
    "enabled": true,
    "manage_volume_offline": true
```

```
  },
  "ikp": {
    "enabled": true,
    "manage_volume_offline": true
  },
  "okm": {
    "enabled": true,
    "manage_volume_offline": true
  },
  "kmip": {
    "enabled": true,
    "manage_volume_offline": true
  }
},
"_links": {
  "self": {
    "href": "/api/security/key-manager-configs"
  }
}
}
----

'''

=== Updating the cluster-level key manager configurations

The following example shows how to modify the
"health_monitor_polling_interval" and "cloud_kms_retry_count" fields.

----

# The API:
PATCH /api/security/key-manager-configs

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-manager-configs' -H
'accept: application/hal+json' -d "{ \"health_monitor_polling_interval\" :
\"20\", \"cloud_kms_retry_count\" : \"5\" }"
----

'''

=== Updating the cluster-level key manager configurations

The following example shows how to modify the "cc_mode" and "passphrase"
fields.
```

```
----

# The API:
PATCH /api/security/key-manager-configs

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-manager-configs' -H
'accept: application/hal+json' -d "{ \"cc_mode_enabled\" : \"true\",
\"passphrase\": \"current_passphrase\" }"
----

'''

=== Shows the keystore level health monitor policy

The following example shows how to retrieve the health monitor policies
for Amazon Web Services and Google Cloud.

----

# The API:
GET /api/security/key-manager-configs

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-manager-
configs?fields=health_monitor_policy.aws,health_monitor_policy.gcp' -H
'accept: application/hal+json'

# The response:
{
"health_monitor_policy": {
  "gcp": {
    "enabled": false,
    "manage_volume_offline": false
  },
  "aws": {
    "enabled": false,
    "manage_volume_offline": false
  }
},
"_links": {
  "self": {
    "href": "/api/security/key-manager-configs"
  }
}
}
```

```
----

'''

=== Updates the keytore level health monitor policy

The following example shows how to modify the Amazon Web Services
"enabled" field and the Google Cloud "manage_volume_offline" field of the
health monitor policy.

----

# The API:
PATCH /api/security/key-manager-configs

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-manager-configs' -H
'accept: application/hal+json' -d '{"health_monitor_policy": {"aws":
{"enabled": "false"}, "gcp": {"manage_volume_offline": "false"}}}'
----

'''




[[ID4f592ea2477bd7b9e12f0cb70a1f3270]]
= Retrieve key manager configurations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-manager-configs`#

*Introduced In:* 9.10

Retrieves key manager configurations.
Retrieves the key manager health monitor policy
(fields=health_monitor_policy).

== Related ONTAP commands

* `security key-manager config show`
* `security key-manager health policy show`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|cc_mode_enabled
|boolean
a|Indicates whether the Common Criteria Mode configuration is enabled.


|cloud_kms_retry_count
|integer
a|Cloud key manager connection retry count. Supported value range of 0-10.


|health_monitor_policy
|link:#health_monitor_policy[health_monitor_policy]
a|Manages the keystore configurations.
```

```
|health_monitor_polling_interval
|integer
a|Health Monitor Polling Period, in minutes. Supported value range of 15-
30 minutes.


|passphrase
|string
a|Current cluster-wide passphrase. This is a required field when setting
the cc_mode_enabled field value to true. This is not audited.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cloud_kms_retry_count": 3,
  "health_monitor_polling_interval": 20,
  "passphrase": "The cluster passphrase of length 64-256 ASCII
characters."
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#akv]
[.api-collapsible-fifth-title]
akv

Azure Key Vault Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#aws]
[.api-collapsible-fifth-title]
aws
```

```
Amazon Web Services Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#barbican]
[.api-collapsible-fifth-title]
barbican

Barbican Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.
```

```
|===


[#gcp]
[.api-collapsible-fifth-title]
gcp

Google Cloud Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#ikp]
[.api-collapsible-fifth-title]
ikp

IBM Key Protect Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.
```

```
|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#kmip]
[.api-collapsible-fifth-title]
kmip

External Key Manager policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#okm]
[.api-collapsible-fifth-title]
okm

Onboard Key Manager policy options


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#health_monitor_policy]
[.api-collapsible-fifth-title]
health_monitor_policy

Manages the keystore configurations.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|akv
|link:#akv[akv]
a|Azure Key Vault Key Management Service policy options


|aws
|link:#aws[aws]
a|Amazon Web Services Key Management Service policy options


|barbican
|link:#barbican[barbican]
a|Barbican Key Management Service policy options


|gcp
```

```
|link:#gcp[gcp]
a|Google Cloud Key Management Service policy options


|ikp
|link:#ikp[ikp]
a|IBM Key Protect Key Management Service policy options


|kmip
|link:#kmip[kmip]
a|External Key Manager policy options


|okm
|link:#okm[okm]
a|Onboard Key Manager policy options


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID628737b57d09cb558596a0d42d173b08]]
= Update key manager configurations

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/key-manager-configs`#

*Introduced In:* 9.10

Updates key manager configurations.
Updates the key manager health monitor policy.
```

```
== Related ONTAP commands

* `security key-manager config modify`
* `security key-manager health policy modify`


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|cc_mode_enabled
|boolean
a|Indicates whether the Common Criteria Mode configuration is enabled.


|cloud_kms_retry_count
|integer
a|Cloud key manager connection retry count. Supported value range of 0-10.


|health_monitor_policy
|link:#health_monitor_policy[health_monitor_policy]
a|Manages the keystore configurations.


|health_monitor_polling_interval
|integer
a|Health Monitor Polling Period, in minutes. Supported value range of 15-
30 minutes.


|passphrase
|string
a|Current cluster-wide passphrase. This is a required field when setting
the cc_mode_enabled field value to true. This is not audited.


|===


.Example request
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "cloud_kms_retry_count": 3,
  "health_monitor_polling_interval": 20,
  "passphrase": "The cluster passphrase of length 64-256 ASCII
characters."
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536139
| Cluster-wide passphrase is incorrect.

| 65536805
| Common Criteria Mode requires an effective cluster version of ONTAP 9.4
or later.

| 65536806
| Passphrase length error.

| 65536807
| MetroCluster cannot be configured while in Common Criteria mode.

| 65536809
| Common Criteria mode is disabled on the cluster. Contact technical
support for assistance in enabling Common Criteria mode.

| 65537301
| The Onboard Key Manager is not enabled.

| 65537302
```

```
| The passphrase field is required when changing cc_mode_enabled to true.

| 65537303
| Modifying polling period requires an effective cluster version of ONTAP
9.10 or later.

| 65537304
| Unable to modify polling period because no external key management is
configured on the cluster.

| 65538404
| Modifying cloud keymanager retry count requires an effective cluster
version of ONTAP 9.11 or later.

| 65539303
| Modifying the health monitor policy requires an effective cluster
version of 9.15 or later.

| 65539304
| The health monitor policy feature is not enabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#akv]
[.api-collapsible-fifth-title]
akv

Azure Key Vault Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#aws]
[.api-collapsible-fifth-title]
aws

Amazon Web Services Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#barbican]
[.api-collapsible-fifth-title]
barbican

Barbican Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#gcp]
[.api-collapsible-fifth-title]
gcp

Google Cloud Key Management Service policy options
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#ikp]
[.api-collapsible-fifth-title]
ikp

IBM Key Protect Key Management Service policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===
```

```
[#kmip]
[.api-collapsible-fifth-title]
kmip

External Key Manager policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#okm]
[.api-collapsible-fifth-title]
okm

Onboard Key Manager policy options


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether health monitor is enabled.


|manage_volume_offline
```

```
|boolean
a|Indicates whether the health monitor manages the volume offline
operation.


|===


[#health_monitor_policy]
[.api-collapsible-fifth-title]
health_monitor_policy

Manages the keystore configurations.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|akv
|link:#akv[akv]
a|Azure Key Vault Key Management Service policy options


|aws
|link:#aws[aws]
a|Amazon Web Services Key Management Service policy options


|barbican
|link:#barbican[barbican]
a|Barbican Key Management Service policy options


|gcp
|link:#gcp[gcp]
a|Google Cloud Key Management Service policy options


|ikp
|link:#ikp[ikp]
a|IBM Key Protect Key Management Service policy options


|kmip
```

```
|link:#kmip[kmip]
a|External Key Manager policy options


|okm
|link:#okm[okm]
a|Onboard Key Manager policy options


|===


[#key_manager_config]
[.api-collapsible-fifth-title]
key_manager_config

Manages the various key manager configuration options.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cc_mode_enabled
|boolean
a|Indicates whether the Common Criteria Mode configuration is enabled.


|cloud_kms_retry_count
|integer
a|Cloud key manager connection retry count. Supported value range of 0-10.


|health_monitor_policy
|link:#health_monitor_policy[health_monitor_policy]
a|Manages the keystore configurations.


|health_monitor_polling_interval
|integer
a|Health Monitor Polling Period, in minutes. Supported value range of 15-
30 minutes.


|passphrase
```

```
|string
a|Current cluster-wide passphrase. This is a required field when setting
the cc_mode_enabled field value to true. This is not audited.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage key managers

:leveloffset: +1


[[IDe20f16e63b08f043abeca6e379044f63]]
= Manage key managers



== Overview

A key manager is a key management solution (software or dedicated
hardware) that enables other ONTAP client modules to securely and
persistently store keys for various uses. For example, WAFL uses the key
management framework to store and retrieve the volume encryption keys that
it uses to encrypt/decrypt data on NVE volumes. A key manager can be
configured at both cluster scope and SVM, with one key manager allowed per
SVM. The key management framework in ONTAP supports two mutually exclusive
modes for persisting keys: external and onboard.

When an SVM is configured with external key management, the keys are
```

```
stored on up to four primary key servers that are external to the system.

Once external key management is enabled for an SVM, primary key servers
can be added or removed using the _/api/security/key-managers/\{uuid}/key-
servers_ endpoint. See [`+POST /security/key-managers/{uuid}/key-
servers+`] and [`+DELETE /security/key-managers/{uuid}/key-
servers/{server}+`] for more details.

Setting up external key management dictates that the required certificates
for securely communicating with the key server are installed prior to
configuring the key manager. To install the required client and server_ca
certificates, use the _/api/security/certificates/_ endpoint.

See [`POST /security/certificates`], [`GET /security/certificates/uuid`]
and [`+DELETE /security/certificates/{uuid}+`] for more details.

When an SVM is configured with the Onboard Key Manager, the keys are
stored in ONTAP in wrapped format using a key hierarchy created using the
salted hash of the passphrase entered when configuring the Onboard Key
Manager. This model fits well for customers who use ONTAP to store their
own data.

== Examples

=== Creating an external key manager with 1 primary key server for a
cluster

The example key manager is configured at the cluster-scope with one
primary key server. Note that the UUIDs of the certificates are those that
are already installed at the cluster-scope. Note the _return_records=true_
query parameter is used to obtain the newly created key manager
configuration.

----

# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=true' -H 'accept: application/hal+json' -d "{
\"external\": { \"client_certificate\": { \"uuid\": \"5fb1701a-d922-11e8-
bfe8-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\":
\"827d7d31-d6c8-11e8-b5bf-005056bb017d\" }],\"servers\": [ { \"server\":
\"10.225.89.33:5696\" } ] } }"

# The response:
```

```
{
"num_records": 1,
"records": [
  {
    "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
    "external": {
      "client_certificate": {
        "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d"
      },
      "server_ca_certificates": [
        {
          "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d"
        }
      ],
      "servers": [
        {
          "server": "10.225.89.33:5696"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-
005056bb017d"
      }
    }
  }
]
}
----

'''

=== Creating an external key manager with two primary key servers

The example key manager is configured at the cluster-scope with two
primary key servers. Note that the UUIDs of the certificates are those
that are already installed at the cluster-scope. Note the
_return_records=true_ query parameter is used to obtain the newly created
key manager configuration.

----

# The API:
POST /api/security/key-managers
```

```
# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=true' -H 'accept: application/hal+json' -d "{
\"external\": { \"client_certificate\": { \"uuid\": \"5fb1701a-d922-11e8-
bfe8-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\":
\"827d7d31-d6c8-11e8-b5bf-005056bb017d\" }],\"servers\": [ { \"server\":
\"104.224.89.33:5696\" }, { \"server\": \"104.224.89.34:5696\" } ] } }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
    "external": {
      "client_certificate": {
        "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d"
      },
      "server_ca_certificates": [
        {
          "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d"
        }
      ],
      "servers": [
        {
          "server": "10.225.89.33:5696"
        },
        {
          "server": "10.225.89.34:5696"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-
005056bb017d"
      }
    }
  }
]
}
----

'''

=== Creating an external key manager with 1 primary key server for an SVM
```

```
The example key manager is configured at the SVM-scope with one primary
key server. Note that the UUIDs of the certificates are those that are
already installed in that SVM. Note the _return_records=true_ query
parameter is used to obtain the newly created key manager configuration.

----

# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=true' -H 'accept: application/hal+json' -d "{
\"svm\": { \"uuid\": \"216e6c26-d6c6-11e8-b5bf-005056bb017d\" },
\"external\": { \"client_certificate\": { \"uuid\": \"91dcaf7c-dbbd-11e8-
9b2c-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\":
\"a4d4b8ba-dbbd-11e8-9b2c-005056bb017d\" }],\"servers\": [ { \"server\":
\"10.225.89.34:5696\" } ] } }"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "80af63f2-dbbf-11e8-9b2c-005056bb017d",
    "svm": {
      "uuid": "216e6c26-d6c6-11e8-b5bf-005056bb017d"
    },
    "external": {
      "client_certificate": {
      "uuid": "91dcaf7c-dbbd-11e8-9b2c-005056bb017d"
      },
      "server_ca_certificates": [
        {
          "uuid": "a4d4b8ba-dbbd-11e8-9b2c-005056bb017d"
        }
      ],
      "servers": [
        {
          "server": "10.225.89.34:5696"
        }
      ]
    },
    "_links": {
      "self": {
```

```
        "href": "/api/security/key-managers/80af63f2-dbbf-11e8-9b2c-
005056bb017d"
      }
    }
  }
]
}
----

'''

=== Creating an onboard key manager for a cluster

The following example shows how to create an onboard key manager for a
cluster with the onboard key manager configured at the cluster-scope.

----

# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers' -H 'accept:
application/hal+json' -d '{ "onboard": { "passphrase": "passphrase" } }'
----

'''

=== Retrieving the key manager configurations for all clusters and SVMs

The following example shows how to retrieve all configured key managers
along with their configurations.

----

# The API:
GET /api/security/key-managers

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers?fields=*' -H
'accept: application/hal+json'

# The response:
{
"records": [
  {
    "configuration": {
```

```
      "uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
      "name" : "default"
    },
    "uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
    "scope": "svm",
    "svm": {
      "uuid": "0f22f8f3-d6c6-11e8-b5bf-005056bb017d",
      "name": "vs0"
    },
    "external": {
      "client_certificate": {
        "uuid": "4cb15482-d6c8-11e8-b5bf-005056bb017d",
        "_links": {
          "self": {
            "href": "/api/security/certificates/4cb15482-d6c8-11e8-b5bf-
005056bb017d/"
          }
        }
      },
      "server_ca_certificates": [
        {
          "uuid": "8a17c858-d6c8-11e8-b5bf-005056bb017d",
          "_links": {
            "self": {
              "href": "/api/security/certificates/8a17c858-d6c8-11e8-b5bf-
005056bb017d/"
            }
          }
        }
      ],
      "servers": [
        {
          "server": "10.2.30.4:5696",
          "timeout": 25,
          "username": "",
          "create_remove_timeout": 10,
          "_links": {
            "self": {
              "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d/key-servers/10.2.30.4:5696/"
            }
          }
        },
        {
          "server": "vs0.local1:3678",
          "timeout": 25,
```

```
          "username": "",
          "secondary_key_servers": "1.1.1.1, secondarykeyserver.com",
          "create_remove_timeout": 10,
          "_links": {
            "self": {
              "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d/key-servers/vs0.local1:3678/"
            }
          }
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d"
      }
    }
  },
  {
   "configuration": {
      "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
      "name" : "default"
    },
    "enabled": true,
    "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
    "scope": "cluster",
    "external": {
      "client_certificate": {
        "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d",
        "_links": {
          "self": {
            "href": "/api/security/certificates/5fb1701a-d922-11e8-bfe8-
005056bb017d/"
          }
        }
      },
      "server_ca_certificates": [
        {
          "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d",
          "_links": {
            "self": {
              "href": "/api/security/certificates/827d7d31-d6c8-11e8-b5bf-
005056bb017d/"
            }
          }
```

```
        }
      ],
      "servers": [
        {
          "server": "10.225.89.33:5696",
          "timeout": 25,
          "username": "",
          "create_remove_timeout": 10,
          "_links": {
            "self": {
              "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-
005056bb017d/key-servers/10.225.89.33:5696/"
            }
          }
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-
005056bb017d"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/key-managers?fields=*"
  }
}
}
----

'''

=== Retrieving the key manager configurations for all clusters and SVMs
(showing Onboard Key Manager)

The following example shows how to retrieve all configured key managers
along with their configurations.

----

# The API:
GET /api/security/key-managers
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers?fields=*' -H
'accept: application/hal+json'

# The response:
{
"records": [
  {
   "configuration": {
      "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
      "name" : "default"
    },
    "enabled": true,
    "uuid": "8ba52e0f-ae22-11e9-b747-005056bb7636",
    "scope": "cluster",
    "onboard": {
      "enabled": true,
      "key_backup": "--------------------------BEGIN
BACKUP--------------------------\n <Backup Data>
\n---------------------------END BACKUP---------------------------\n"
    },
    "volume_encryption": {
      "supported": false,
      "message": "The following nodes do not support volume granular
encryption: ntap-vsim2.",
      "code": 65536935
    },
    "is_default_data_at_rest_encryption_disabled": false
  }
],
"num_records": 1
}
----

'''

=== Retrieving expensive fields such as, status.code and status.message,
associated with a key manager.

These values are not retrieved by default with the 'fields=*' option.
The following example shows how to retrieve the expensive objects
associated with a key manager.

----

# The API:
```

```
GET /api/security/key-managers

# The call:
curl -X GET "https://<mgmt-ip>/api/security/key-
managers?fields=status.message,status.code" -H 'accept:
application/hal+jon'

# The response:
{
"records": [
  {
    "configuration": {
      "uuid": "ac305d46-aef4-11e9-ad3c-005056bb7636",
      "name" : "default"
    },
    "enabled": true,
    "uuid": "ac305d46-aef4-11e9-ad3c-005056bb7636",
    "status": {
      "message": "No action needed at this time.",
      "code": 65537200
    },
    "_links": {
      "self": {
        "href": "/api/security/key-managers/ac305d46-aef4-11e9-ad3c-
005056bb7636"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/key-managers?fields=status.message,status.code"
  }
}
}
----

'''

=== Retrieving a specific key manager configuration

The following example shows how to retrieve a specific key manager
configuration.

----
```

```
# The API:
GET /api/security/key-managers/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>?fields=*'
-H 'accept: application/hal+json'

# The response:
{
"configuration": {
      "uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
      "name" : "default"
},
"enabled": true,
"uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
"scope": "svm",
"svm": {
  "uuid": "0f22f8f3-d6c6-11e8-b5bf-005056bb017d",
  "name": "vs0"
},
"external": {
  "client_certificate": {
    "uuid": "4cb15482-d6c8-11e8-b5bf-005056bb017d",
    "_links": {
      "self": {
        "href": "/api/security/certificates/4cb15482-d6c8-11e8-b5bf-
005056bb017d/"
      }
    }
  },
  "server_ca_certificates": [
    {
      "uuid": "8a17c858-d6c8-11e8-b5bf-005056bb017d",
      "_links": {
        "self": {
          "href": "/api/security/certificates/8a17c858-d6c8-11e8-b5bf-
005056bb017d/"
        }
      }
    }
  ],
  "servers": [
    {
      "server": "10.2.30.4:5696",
      "timeout": 25,
      "username": "",
```

```
      "create_remove_timeout": 10,
      "_links": {
        "self": {
          "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d/key-servers/10.2.30.4:5696/"
        }
      }
    },
    {
      "server": "vs0.local1:3678",
      "timeout": 25,
      "username": "",
      "create_remove_timeout": 10,
      "_links": {
        "self": {
          "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d/key-servers/vs0.local1:3678/"
        }
      }
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-
005056bb017d"
  }
}
}
----

'''

=== Updating the configuration of an external key manager

The following example shows how to update the server_ca configuration of
an external key manager.

----

# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{ \"external\": {
```

```
\"server_ca_certificates\": [ { \"uuid\": \"23b05c58-d790-11e8-b5bf-
005056bb017d\" }] } }"
----

'''

=== Updating the passphrase of an Onboard Key Manager

The following example shows how to update the passphrase of a given key
manager.

----

# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{ \"onboard\": {
\"existing_passphrase\": \"existing_passphrase\", \"passphrase\":
\"new_passphrase\" } }"
----

'''

=== Synchronizing the passphrase of the Onboard Key Manager on a cluster

The following example shows how to synchronize the passphrase on a cluster
where the Onboard Key Manager is already configured.

----

# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{  \"onboard\": {
\"existing_passphrase\": \"existing_passphrase\",    \"synchronize\": true
}}"
----

'''

=== Configuring the Onboard Key Manager on a cluster

The following example shows how to configure the Onboard Key Manager on a
```

```
cluster where the Onboard Key Manager is not configured, but is configured
on an MetroCluster partner cluster.

----

# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=false' -H  'accept: application/hal+json' -H
"Content-Type: application/json" -d "{  \"onboard\": {    \"passphrase\":
\"passphrase\",    \"synchronize\": true  }}"
----

'''

=== Deleting a configured key manager

The following example shows how to delete a key manager given its UUID.

----

# The API:
DELETE /api/security/key-managers/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json'
----

'''

=== Adding a primary key server to an external key manager

The following example shows how to add a primary key server to an external
key manager.

----

# The API:
POST /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers?return_records=true' -H 'accept: application/hal+json' -d "{
\"server\": \"10.225.89.34:5696\" }"
```

```
# The response:
{
"num_records": 1,
"records": [
  {
    "server": "10.225.89.34:5696",
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers/10.225.89.34%3A5696"
      }
    }
  }
]
}
----

'''

=== Adding 2 primary key servers to an external key manager

The following example shows how to add 2 primary key servers to an
external key manager. Note that the _records_ property is used to add
multiple primary key servers to the key manager in a single API call.

----

# The API:
POST /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers?return_records=true' -H 'accept: application/hal+json' -d "{
\"records\": [ { \"server\": \"10.225.89.34:5696\" }, { \"server\":
\"10.225.89.33:5696\" } ] }"

# The response:
{
"num_records": 1,
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers/"
```

```
      }
    }
  }
]
}
----

'''

=== Retrieving all the key servers configured in an external key manager

The following example shows how to retrieve all key servers configured in
an external key manager.

----

# The API:
GET /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers?fields=*' -H 'accept: application/hal+json'

# The response:
{
"records": [
  {
    "uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
    "server": "10.225.89.33:5696",
    "timeout": 25,
    "username": "",
    "secondary_key_servers": [
      "1.1.1.1",
      "secondarykeyserver.com"
    ],
    "create_remove_timeout": 10,
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers/10.225.89.33%3A5696"
      }
    }
  },
  {
    "uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
    "server": "10.225.89.34:5696",
```

```
    "timeout": 25,
    "username": "",
    "create_remove_timeout": 10,
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers/10.225.89.34%3A5696"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers?fields=*"
  }
}
}
----

'''

=== Retrieving a specific primary key server (and any associated secondary
key servers) configured in an external key manager

The following example shows how to retrieve a specific primary key server
(and any associated secondary key servers) configured in an external key
manager.

----

# The API:
GET /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}?fields=*' -H 'accept: application/hal+json'

# The response:
{
"uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
"server": "10.225.89.34:5696",
"timeout": 25,
"username": "",
"secondary_key_servers": [
```

```
  "1.1.1.1",
  "secondarykeyserver.com"
],
"create_remove_timeout": 10,
"_links": {
  "self": {
    "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-
005056bb017d/key-servers/10.225.89.34:5696"
  }
}
}
----

'''

=== Retrieving a specific primary key server (and any associated secondary
key servers) (and connectivity, an expensive field) configured in an
external key manager

The following example shows how to retrieve a specific primary key server
(and any associated secondary key servers) configured in an external key
manager.

----

# The API:
GET /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}?fields=**' -H 'accept: application/hal+json'

# The response:
{
"uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
"server": "10.225.89.34:5696",
"timeout": 25,
"username": "",
"secondary_key_servers": [
  "1.1.1.1",
  "secondarykeyserver.com"
],
"create_remove_timeout": 10,
"connectivity": {
  "cluster_availability": true,
  "node_states": [
```

```
    {
      "node": {
        "name": "sti65-vsim-ucs148i",
        "uuid": "661843b3-a0e5-11ed-81ef-005056a7306b"
      },
      "state": "available"
    },
    {
      "node": {
        "name": "sti65-vsim-ucs148j",
        "uuid": "551843b3-a0e5-11ed-81ef-005056a7306b"
      },
      "state": "not_responding"
    }
  ]
}
}
----

'''

=== Retrieving the connectivity status of a specific node for a specific
primary key server configured in an external key manager

The following example shows how to retrieve the connectivity status for a
specific node for a specific primary key server configured in an external
key manager.

----

# The API:
GET /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/43e0c191-dc5c-
11e8-9b2c-005056bb017d/key-
servers/10.225.89.34:5696?fields=connectivity&connectivity.node_states.nod
e.name=sti65-vsim-ucs148i&return_unmatched_nested_array_objects=false' -H
'accept: application/hal+json'

# The response:
{
"uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
"server": "10.225.89.34:5696",
"connectivity": {
  "cluster_availability": true,
```

```
  "node_states": [
    {
      "node": {
        "name": "sti65-vsim-ucs148i",
        "uuid": "661843b3-a0e5-11ed-81ef-005056a7306b"
      },
      "state": "available"
    }
  ]
}
}
----

'''

=== Updating a specific primary key server configuration configured in an
external key manager

The following example shows how to update a specific primary key server
configured in an external key manager.

----

# The API:
PATCH /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}' -H 'accept: application/hal+json' -d "{ \"timeout\": 45
}"
----

'''

=== When the 'secondary_key_servers' field is populated in the PATCH API,
the list of secondary key servers

=== associated with the primary key servers is replaced by the list of
secondary key servers specified in the

=== 'secondary_key_servers' field.

The following example shows how to update the set of secondary key servers
associated with a primary key server.

----
```

```
# The API:
PATCH /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}' -H 'accept: application/hal+json' -d "{
\"secondary_key_servers\": [ \"1.1.1.1\", \"secondarykeyserver.com\" ] }"
----

'''

=== Deleting a primary key server from an external key manager

The following example shows how to delete a primary key server from an
external key manager.

----

# The API:
DELETE /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}' -H 'accept: application/hal+json'
----

'''

=== Bypass the out of quorum checks when deleting a primary key server
from an external key manager

The following example shows how to bypass the out of quorum checks when
deleting a primary key server from an external key manager.

----

# The API:
DELETE /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-
servers/{server}?force=true' -H 'accept: application/hal+json'
----

'''
```

```
[[ID189d317580e98f4d21f3b82f02b56174]]
= Retrieve key managers

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-managers`#

*Introduced In:* 9.6

Retrieves key managers.

== Expensive properties

There is an added computational cost to retrieving these properties. They
are excluded from default GET results and must be explicitly requested
using the `fields` query parameter. For more details, see
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields]. Additionally, these fields
are unavailable for inactive configurations as they are only relevant to
active configurations.

* `status.message`
* `status.code`

== Examples

* To retrieve basic information about a key manager:

----
GET /security/key-managers
----

* To retrieve specific fields, including expensive properties:

----
GET /security/key-managers/?fields=status.code,status.message
----

== Related ONTAP commands

* `security key-manager show-key-store`
* `security key-manager external show`
* `security key-manager external show-status`
* `security key-manager onboard show-backup`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|configuration.uuid
|string
|query
|False
a|Filter by configuration.uuid

* Introduced in: 9.16


|configuration.name
|string
|query
|False
a|Filter by configuration.name

* Introduced in: 9.16


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|policy
|string
|query
```

```
|False
a|Filter by policy

* Introduced in: 9.9


|status.code
|integer
|query
|False
a|Filter by status.code

* Introduced in: 9.7


|status.message
|string
|query
|False
a|Filter by status.message

* Introduced in: 9.7


|onboard.enabled
|boolean
|query
|False
a|Filter by onboard.enabled


|onboard.key_backup
|string
|query
|False
a|Filter by onboard.key_backup

* Introduced in: 9.7


|enabled
|boolean
|query
|False
a|Filter by enabled

* Introduced in: 9.16
```

```
|is_default_data_at_rest_encryption_disabled
|boolean
|query
|False
a|Filter by is_default_data_at_rest_encryption_disabled

* Introduced in: 9.7


|volume_encryption.message
|string
|query
|False
a|Filter by volume_encryption.message

* Introduced in: 9.7


|volume_encryption.supported
|boolean
|query
|False
a|Filter by volume_encryption.supported

* Introduced in: 9.7


|volume_encryption.code
|integer
|query
|False
a|Filter by volume_encryption.code

* Introduced in: 9.7


|external.client_certificate.name
|string
|query
|False
a|Filter by external.client_certificate.name

* Introduced in: 9.8
```

```
|external.client_certificate.uuid
|string
|query
|False
a|Filter by external.client_certificate.uuid


|external.servers.connectivity.node_states.node.name
|string
|query
|False
a|Filter by external.servers.connectivity.node_states.node.name

* Introduced in: 9.13


|external.servers.connectivity.node_states.node.uuid
|string
|query
|False
a|Filter by external.servers.connectivity.node_states.node.uuid

* Introduced in: 9.13


|external.servers.connectivity.node_states.state
|string
|query
|False
a|Filter by external.servers.connectivity.node_states.state

* Introduced in: 9.13


|external.servers.connectivity.cluster_availability
|boolean
|query
|False
a|Filter by external.servers.connectivity.cluster_availability

* Introduced in: 9.7


|external.servers.username
|string
|query
|False
```

```
a|Filter by external.servers.username


|external.servers.server
|string
|query
|False
a|Filter by external.servers.server


|external.servers.timeout
|integer
|query
|False
a|Filter by external.servers.timeout

* Max value: 60
* Min value: 1


|external.servers.secondary_key_servers
|string
|query
|False
a|Filter by external.servers.secondary_key_servers

* Introduced in: 9.8


|external.server_ca_certificates.name
|string
|query
|False
a|Filter by external.server_ca_certificates.name

* Introduced in: 9.8


|external.server_ca_certificates.uuid
|string
|query
|False
a|Filter by external.server_ca_certificates.uuid


|uuid
|string
```

```
|query
|False
a|Filter by uuid


|scope
|string
|query
|False
a|Filter by scope


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.
```

```
* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#security_key_manager[security_key_manager]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
```

```
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "configuration": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      },
      "external": {
        "client_certificate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "server_ca_certificates": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "string",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "servers": [
          {
            "_links": {
```

```
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "connectivity": {
              "node_states": [
                {
                  "node": {
                    "_links": {
                      "self": {
                        "href": "/api/resourcelink"
                      }
                    },
                    "name": "node1",
                    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
                  },
                  "state": "not_responding"
                }
              ]
            },
            "secondary_key_servers": "secondary1.com, 10.2.3.4",
            "server": "keyserver1.com:5698",
            "timeout": 60,
            "username": "admin"
          }
        ]
      },
      "onboard": {
        "existing_passphrase": "The cluster password of length 32-256
ASCII characters.",
        "key_backup": "'--------------------------BEGIN
BACKUP--------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAAAxBFWWAAAAACEAAAAAAAAA
QAAAAAAAAABzDyyVAAAAALI5Jsjvy6gUxnT78KoDKXHYb6sSeraM00quOULY6BeV
n6dMFxuErCD1lbERaOQZSuaYy1p8oQHtTEfGMLZM4TYiAAAAAAAAACgAAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAAA
IgAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/LRzU
QRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAAADV1Vd/AAAAAMFM9Q229Bhp
mDaTSdqku5DCd8wG+fOZSr4bx4JT5WHvV/r5gJnXDQQAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJsb2IA
AQAAAAMAAAAYAQAAAAAALgePkcAAAAAIgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAA
AAAAAAAAAAACAAAAAAAJAGr3tJA/LRzUQRHwv+1aWvAAAAAAAAAAACIAAAAAAAAA
KAAAAAAAAACIlCHZAAAAAAAAAAAAAAAAAAgAAAAAAQCafcabsxRXMM7gxhLRrzxh
AAAAAAAAAAAkAAAAAAAAAIAAAAAAAAAA2JjQBQAAAACt4IqXcNpVggahl0axLsN4
yQjnNVKWY7mANB29O42hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAIAAAAAAAkA
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAAIgAAAAAAAAoAAAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAAACQAAAAAAAA
gAAAAAAAAACnMmUtAAAAAGVk8AtPzENFgsGdsFvnmucmYrlQCsFew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
----------END BACKUP----------'",
        "synchronize": true
      },
      "policy": "string",
      "scope": "string",
      "status": {
        "code": 346758,
        "message": "This cluster is part of a MetroCluster configuration.
Use the REST API POST method security/key_managers/ with the synchronize
option and the same passphrase on the partner cluster before proceeding
with any key manager operations.  Failure to do so could lead to
switchover or switchback failure."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string",
      "volume_encryption": {
        "code": 346758,
        "message": "No platform support for volume encryption in following
nodes - node1, node2."
      }
    }
  ]
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
```

```
|link:#href[href]
a|

|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#client_certificate]
[.api-collapsible-fifth-title]
client_certificate

Client certificate (name and UUID)


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#server_ca_certificates]
[.api-collapsible-fifth-title]
server_ca_certificates

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===
```

```
[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state
```

```
The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.
```

```
|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#key_server_readcreate]
[.api-collapsible-fifth-title]
key_server_readcreate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|secondary_key_servers
|string
a|A comma delimited string of the secondary key servers associated with
the primary key server.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used.


|timeout
|integer
```

```
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|Username credentials for connecting with the key server.


|===


[#external]
[.api-collapsible-fifth-title]
external

Configures external key management


[cols=3*,options=header]
|===
|Name
|Type
|Description

|client_certificate
|link:#client_certificate[client_certificate]
a|Client certificate (name and UUID)


|server_ca_certificates
|array[link:#server_ca_certificates[server_ca_certificates]]
a|The array of certificates that are common for all the keyservers per
SVM.


|servers
|array[link:#key_server_readcreate[key_server_readcreate]]
a|The set of external key servers.


|===


[#onboard]
[.api-collapsible-fifth-title]
onboard
```

```
Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Is the onboard key manager enabled?


|existing_passphrase
|string
a|The cluster-wide passphrase. This is not audited.


|key_backup
|string
a|Backup of the onboard key manager's key hierarchy. It is required to
save this backup after configuring the onboard key manager to help in the
recovery of the cluster in case of catastrophic failures.


|===


[#status]
[.api-collapsible-fifth-title]
status

Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
```

```
a|Code corresponding to the status message. Returns 0 if the setup is
complete. This is an advanced property; there is an added computational
cost to retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|message
|string
a|Current state of the key manager indicating any additional steps to
perform to finish the setup. This is an advanced property; there is an
added computational cost to retrieving its value. The property is not
populated for either a collection GET or an instance GET unless it is
explicitly requested using the `fields` query parameter or GET for all
advanced properties is enabled.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#volume_encryption]
[.api-collapsible-fifth-title]
volume_encryption

Indicates whether volume encryption is supported in the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns a 0 if volume
encryption is supported in all nodes of the cluster.


|message
|string
a|Reason for not supporting volume encryption.


|supported
|boolean
a|Set to true when volume encryption support is available on all nodes of
the cluster.


|===


[#security_key_manager]
[.api-collapsible-fifth-title]
security_key_manager

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|is_default_data_at_rest_encryption_disabled
|boolean
a|Indicates whether default data-at-rest encryption is disabled in the
cluster. This field is deprecated in ONTAP 9.8 and later. Use the
"software_data_encryption.disabled_by_default" of /api/security endpoint.

* Default value:
* Introduced in: 9.7
* x-ntap-readModify: true
* x-nullable: true


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


|policy
|string
a|Security policy associated with the key manager. This value is currently
ignored if specified for the onboard key manager.


|scope
|string
```

```
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|status
|link:#status[status]
a|Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|volume_encryption
|link:#volume_encryption[volume_encryption]
a|Indicates whether volume encryption is supported in the cluster.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID24e2178b8dd555789050d706f7b2dae8]]
= Create a key manager

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/key-managers`#
```

```
*Introduced In:* 9.6

Creates a key manager.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create a key
manager.
* `external.client_certificate` - Client certificate. Required only when
creating an external key manager.
* `external.server_ca_certificates` - Server CA certificates. Required
only when creating an external key manager.
* `external.servers.server` - Primary Key servers. Required only when
creating an external key manager.
* `onboard.passphrase` - Cluster-wide passphrase. Required only when
creating an Onboard Key Manager.
* `synchronize` - Synchronizes missing onboard keys on any node in the
cluster. Required only when creating an Onboard Key Manager at the partner
site of a MetroCluster configuration.

== Optional parameters:

* `create_inactive` - Create a configuration without enabling it. This
flag is set to "false" by default.

== Related ONTAP commands

* `security key-manager external enable`
* `security key-manager onboard enable`
* `security key-manager onboard sync`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|create_inactive
|boolean
|query
|False
```

```
a|Indicates whether to create an active or inactive configuration.

* Introduced in: 9.16


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.
```

```
|enabled
|boolean
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


|policy
|string
a|Security policy associated with the key manager. This value is currently
ignored if specified for the onboard key manager.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "configuration": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "external": {
    "client_certificate": {
      "name": "string",
```

```
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": [
      {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "servers": [
      {
        "secondary_key_servers": "secondary1.com, 10.2.3.4",
        "server": "keyserver1.com:5698",
        "timeout": 60,
        "username": "admin"
      }
    ]
  },
  "is_default_data_at_rest_encryption_disabled": true,
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII
characters.",
    "key_backup": "'--------------------------BEGIN
BACKUP--------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAAAxBFWWAAAAACEAAAAAAAAA
QAAAAAAAAABzDyyVAAAAALI5Jsjvy6gUxnT78KoDKXHYb6sSeraM00quOULY6BeV
n6dMFxuErCD1lbERaOQZSuaYy1p8oQHtTEfGMLZM4TYiAAAAAAAAACgAAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAAA
IgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/LRzU
QRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAAADV1Vd/AAAAAMFM9Q229Bhp
mDaTSdqku5DCd8wG+fOZSr4bx4JT5WHvV/r5gJnXDQQAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJsb2IA
AQAAAAMAAAAYAQAAAAAALgePkcAAAAAIgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAA
AAAAAAAAAAACAAAAAAAJAGr3tJA/LRzUQRHwv+1aWvAAAAAAAAAAACIAAAAAAAAA
KAAAAAAAAACIlCHZAAAAAAAAAAAAAAAAAgAAAAAAAQCafcabsxRXMM7gxhLRrzxh
AAAAAAAAAAAkAAAAAAAAAIAAAAAAAAAA2JjQBQAAAACt4IqXcNpVggahl0axLsN4
yQjnNVKWY7mANB29O42hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAIAAAAAAAkA
```

```
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAAIgAAAAAAAAAoAAAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAAACQAAAAAAAA
gAAAAAAAACnMmUtAAAAGVk8AtPzENFgsGdsFvnmucmYrlQCsFew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
----------END BACKUP----------'",
    "passphrase": "The cluster password of length 32-256 ASCII
characters."
  },
  "policy": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
====

== Response
```


Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262224
| Failed to contact the peer cluster.

| 262228
| Failed to contact the peer cluster.

| 65536038
| A maximum of 4 active primary key servers are allowed.

| 65536214
| Failed to generate cluster key encryption key.

| 65536216
| Failed to add cluster key encryption key.

| 65536310
| Failed to setup the Onboard Key Manager because the MetroCluster peer is
unhealthy.

| 65536341
| Failed to setup the Onboard Key Manager because the MetroCluster peer is
unhealthy.

| 65536508
| The platform does not support data at rest encryption.

| 65536821
| The certificate is not installed.

| 65536823
| The SVM has key manager already configured.

| 65536824
| Multitenant key management is not supported in MetroCluster
configurations.
```

```
| 65536834
| Failed to get existing key-server details for the SVM.

| 65536852
| Failed to query supported KMIP protocol versions.

| 65536870
| Key management servers already configured.

| 65536871
| Duplicate key management servers exist.

| 65536876
| External key management requires client and server CA certificates
installed and with one or more key servers provided.

| 65536878
| External key management cannot be configured as one or more volume
encryption keys of the SVM are stored in cluster key management server.

| 65536895
| External key manager cannot be configured because this cluster is part
of a MetroCluster configuration and the partner site of this MetroCluster
configuration has Onboard Key Manager configured.

| 65536900
| The Onboard Key Manager cannot be configured because this cluster is
part of a MetroCluster configuration and the partner site has the external
key manager configured.

| 65536903
| The Onboard Key Manager has failed to configure on some nodes in the
cluster. Use the CLI to sync the Onboard Key Manager configuration on
failed nodes.

| 65536906
| The Onboard Key Manager has already been configured at the partner site.
Use the CLI to sync the Onboard Key Manager with the same passphrase.

| 65536916
| The Onboard Key Manager is only supported for an admin SVM.

| 65536920
| The Onboard Key Manager passphrase length is incorrect.

| 65537240
| The Onboard Key Manager passphrase must be provided when performing a
```

```
POST/synchronize operation.

| 65537241
| The Onboard Key Manager existing_passphrase must not be provided when
performing a POST/synchronize operation.

| 65537244
| Unable to sync/create Onboard Key Manager on the local cluster; Onboard
Key Manager is already configured on the cluster.

| 65537245
| Unable to sync/create Onboard Key Manager on the local cluster; Onboard
Key Manager is not configured on the partner cluster.

| 65537246
| Unable to sync/create Onboard Key Manager on local cluster. This cluster
is not part of a MetroCluster configuration.

| 65537247
| Internal error. Unable to sync the Onboard Key Manager on local cluster.

| 65537248
| Unable to sync the Onboard Key Manager on local cluster.

| 65538111
| The key manager policy is invalid.

| 65538120
| The key manager policy is not supported on the admin SVM.

| 65539216
| The Admin SVM has a key manager already configured.

| 65539221
| Failed to configure the Onboard Key Manager because the MetroCluster
peer cluster is unhealthy. Verify that the peer cluster is online and
healthy.

| 65539500
| Cannot create an inactive external key manager on this SVM because
inactive external key managers can only be created on the admin SVM.

| 65539501
| Cannot create an inactive external key manager on the admin SVM because
an external key manager already exists on the admin SVM.

| 65539503
```

```
| Cannot create an inactive external key manager on the admin SVM while
MetroCluster is configured.

| 65539504
| An effective cluster version of ONTAP 9.16.1 or later is required to
create an inactive external key manager on the admin SVM.

| 65539511
| Cannot create an inactive Onboard Key Manager on the admin SVM while
MetroCluster is configured.

| 65539512
| An effective cluster version of ONTAP 9.16.1 or later is required to
create an inactive Onboard Key Manager configuration on the admin SVM.

| 65539580
| Failed to create inactive Onboard Key Manager configuration.

| 65539581
| Cannot create an inactive Onboard Key Manager on the admin SVM because
an inactive Onboard Key Manager configuration already exists on the admin
SVM.

| 65539582
| Cannot specify the configuration name parameter. Only one Onboard Key
Manager is supported for the admin SVM.

| 65539704
| The key manager cannot be configured because the SVM has NAE volumes.

| 65539706
| The Onboard Key Manager is already configured. Use the PATCH to sync all
nodes with the Onboard Key Manager configuration.

| 66060338
| Failed to establish secure connection for a key management server due to
incorrect server_ca certificates.

| 66060339
| Failed to establish secure connection for a key management server due to
incorrect client certificates.

| 66060340
| Failed to establish secure connection for a key management server due to
Cryptsoft error.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
```

```
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#client_certificate]
[.api-collapsible-fifth-title]
client_certificate

Client certificate (name and UUID)


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#server_ca_certificates]
[.api-collapsible-fifth-title]
server_ca_certificates

Security certificate object reference


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.
```

```
|===


[#key_server_readcreate]
[.api-collapsible-fifth-title]
key_server_readcreate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|secondary_key_servers
|string
a|A comma delimited string of the secondary key servers associated with
the primary key server.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|Username credentials for connecting with the key server.


|===


[#external]
[.api-collapsible-fifth-title]
external

Configures external key management


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|client_certificate
|link:#client_certificate[client_certificate]
a|Client certificate (name and UUID)


|server_ca_certificates
|array[link:#server_ca_certificates[server_ca_certificates]]
a|The array of certificates that are common for all the keyservers per
SVM.


|servers
|array[link:#key_server_readcreate[key_server_readcreate]]
a|The set of external key servers.


|===


[#onboard]
[.api-collapsible-fifth-title]
onboard

Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Is the onboard key manager enabled?


|existing_passphrase
|string
a|The cluster-wide passphrase. This is not audited.
```

```
|key_backup
|string
a|Backup of the onboard key manager's key hierarchy. It is required to
save this backup after configuring the onboard key manager to help in the
recovery of the cluster in case of catastrophic failures.


|passphrase
|string
a|The cluster-wide passphrase. This is not audited.


|synchronize
|boolean
a|Synchronizes missing onboard keys on any node in the cluster. If a node
is added to a cluster that has onboard key management configured, the
synchronize operation needs to be performed in a PATCH operation. In a
MetroCluster configuration, if onboard key management is enabled on one
site, then the synchronize operation needs to be run as a POST operation
on the remote site providing the same passphrase.


|===


[#status]
[.api-collapsible-fifth-title]
status

Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns 0 if the setup is
complete. This is an advanced property; there is an added computational
cost to retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
```

```
the `fields` query parameter or GET for all advanced properties is
enabled.


|message
|string
a|Current state of the key manager indicating any additional steps to
perform to finish the setup. This is an advanced property; there is an
added computational cost to retrieving its value. The property is not
populated for either a collection GET or an instance GET unless it is
explicitly requested using the `fields` query parameter or GET for all
advanced properties is enabled.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume_encryption]
[.api-collapsible-fifth-title]
volume_encryption
```

```
Indicates whether volume encryption is supported in the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns a 0 if volume
encryption is supported in all nodes of the cluster.


|message
|string
a|Reason for not supporting volume encryption.


|supported
|boolean
a|Set to true when volume encryption support is available on all nodes of
the cluster.


|===


[#security_key_manager]
[.api-collapsible-fifth-title]
security_key_manager

[cols=3*,options=header]
|===
|Name
|Type
|Description

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
```

```
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


|policy
|string
a|Security policy associated with the key manager. This value is currently
ignored if specified for the onboard key manager.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd302d57c465e9a8bff3d8f75347f14ea]]
= Delete key managers

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/key-managers/{uuid}`#

*Introduced In:* 9.6

Deletes a key manager.

== Related ONTAP commands

* `security key-manager external disable`
* `security key-manager onboard disable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
```

```
|string
|path
|True
a|Key manager UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|65536208
//end row
//start row
|Failed to delete the SVM Key ID.
//end row
//start row
|
//end row
//start row
|65536233
//end row
//start row
|Internal error. Deletion of km_wrapped_kdb key database has failed for
the Onboard Key Manager.
//end row
//start row
|
//end row
//start row
|65536234
//end row
```

```
//start row
|Internal error. Deletion of cluster_kdb key database has failed for the
Onboard Key Manager.
//end row
//start row
|
//end row
//start row
|65536239
//end row
//start row
|Encrypted volumes are found for the SVM.
//end row
//start row
|
//end row
//start row
|65536242
//end row
//start row
|One or more self-encrypting drives are assigned an authentication key.
//end row
//start row
|
//end row
//start row
|65536243
//end row
//start row
|Cannot determine authentication key presence on one or more self-
encrypting drives.
//end row
//start row
|
//end row
//start row
|65536800
//end row
//start row
|Failed to lookup onboard keys.
//end row
//start row
|
//end row
//start row
|65536813
```

```
//end row
//start row
|Encrypted kernel core files found.
//end row
//start row
|
//end row
//start row
|65536817
//end row
//start row
|Failed to determine if key manager is safe to disable.
//end row
//start row
|
//end row
//start row
|65536827
//end row
//start row
|Failed to determine if the SVM has any encrypted volumes.
//end row
//start row
|
//end row
//start row
|65536828
//end row
//start row
|External key management is not enabled for the SVM.
//end row
//start row
|
//end row
//start row
|65536867
//end row
//start row
|Encrypted volumes are found for the SVM.
//end row
//start row
|
//end row
//start row
|196608301
//end row
```

```
//start row
|Failed to determine the type of encryption.
//end row
//start row
|
//end row
//start row
|196608305
//end row
//start row
|NAE aggregates are found in the cluster.
//end row
//start row
|Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.

* name: KEYMANAGER_MESSAGE_ERR_KM_DISABLE_ENC_CORE_CHECK_TIMEOUT
message: Failed to disable the key manager because of a timeout when
checking for encrypted cores.
//end row
|===
//end table



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb41e88624818cc0689504447aec8ac88]]
= Retrieve key managers

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-managers/{uuid}`#

*Introduced In:* 9.6

Retrieves key managers.

== Expensive properties

There is an added computational cost to retrieving these properties. They
are excluded from default GET results and must be explicitly requested
```

```
using the `fields` query parameter. For more details, see
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields]. Additionally, these fields
are unavailable for inactive configurations as they are only relevant to
active configurations.

* `status.message`
* `status.code`

== Examples

* To retrieve basic information about a key manager:

----
GET /security/key-managers/{uuid}
----

* To retrieve specific fields, including expensive properties:

----
GET /security/key-managers/{uuid}?fields=status.message
----

== Related ONTAP commands

* `security key-manager show-key-store`
* `security key-manager external show`
* `security key-manager external show-status`
* `security key-manager onboard show-backup`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Key manager UUID
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|is_default_data_at_rest_encryption_disabled
|boolean
a|Indicates whether default data-at-rest encryption is disabled in the
cluster. This field is deprecated in ONTAP 9.8 and later. Use the
"software_data_encryption.disabled_by_default" of /api/security endpoint.

* Default value: 1
```

```
* Introduced in: 9.7
* x-ntap-readModify: true
* x-nullable: true


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


|policy
|string
a|Security policy associated with the key manager. This value is currently
ignored if specified for the onboard key manager.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|status
|link:#status[status]
a|Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|volume_encryption
|link:#volume_encryption[volume_encryption]
a|Indicates whether volume encryption is supported in the cluster.


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configuration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "external": {
    "client_certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "servers": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "connectivity": {
          "node_states": [
            {
              "node": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "node1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              },
              "state": "not_responding"
            }
          ]
        },
        "secondary_key_servers": "secondary1.com, 10.2.3.4",
        "server": "keyserver1.com:5698",
        "timeout": 60,
        "username": "admin"
      }
    ]
  },
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII
characters.",
    "key_backup": "'--------------------------BEGIN
BACKUP--------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAAAxBFWWAAAAACEAAAAAAAAA
QAAAAAAAAABzDyyVAAAAALI5Jsjvy6gUxnT78KoDKXHYb6sSeraM00quOULY6BeV
n6dMFxuErCD1lbERaOQZSuaYy1p8oQHtTEfGMLZM4TYiAAAAAAAAACgAAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAAA
IgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/LRzU
QRHwv+1aWvAAAAAAAAAAACQAAAAAAAAgAAAAAAAAADV1Vd/AAAAAMFM9Q229Bhp
mDaTSdqku5DCd8wG+fOZSr4bx4JT5WHvV/r5gJnXDQQAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

```
KAAAAAAAAACIlCHZAAAAAAAAAAAAAAAAAgAAAAAAQCafcabsxRXMM7gxhLRrzxh
AAAAAAAAAAAkAAAAAAAAAIAAAAAAAAAA2JjQBQAAAACt4IqXcNpVggahl0axLsN4
yQjnNVKWY7mANB29O42hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAIAAAAAAAkA
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAIgAAAAAAAAoAAAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAAACQAAAAAAAA
gAAAAAAAACnMmUtAAAAAGVk8AtPzENFgsGdsFvnmucmYrlQCsFew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
----------END BACKUP----------'",
    "synchronize": true
  },
  "policy": "string",
  "scope": "string",
  "status": {
    "code": 346758,
    "message": "This cluster is part of a MetroCluster configuration. Use
the REST API POST method security/key_managers/ with the synchronize
option and the same passphrase on the partner cluster before proceeding
with any key manager operations.  Failure to do so could lead to
switchover or switchback failure."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string",
  "volume_encryption": {
    "code": 346758,
    "message": "No platform support for volume encryption in following
nodes - node1, node2."
  }
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536894
| This cluster is part of a MetroCluster configuration. Configure an
external key manager on the partner cluster providing the same key servers
before proceeding with any key manager operations.

| 65537201
| There are no key servers configured for this SVM in the local cluster.

| 65537202
| There are no key servers configured for this SVM in the remote cluster.

| 65537203
| Internal error. Failed to check for key servers on partner cluster.

| 65537204
| This cluster is part of a MetroCluster configuration. Configure an
external key manager on the partner cluster providing the same key servers
before proceeding with any key manager operations.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#client_certificate]
[.api-collapsible-fifth-title]
client_certificate

Client certificate (name and UUID)
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#server_ca_certificates]
[.api-collapsible-fifth-title]
server_ca_certificates

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Certificate name


|uuid
```

```
|string
a|Certificate UUID


|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
```

```
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#key_server_readcreate]
[.api-collapsible-fifth-title]
key_server_readcreate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|secondary_key_servers
|string
a|A comma delimited string of the secondary key servers associated with
the primary key server.


|server
|string
a|External key server for key management. If no port is provided, a
```

```
default port of 5696 is used.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|Username credentials for connecting with the key server.


|===


[#external]
[.api-collapsible-fifth-title]
external

Configures external key management


[cols=3*,options=header]
|===
|Name
|Type
|Description

|client_certificate
|link:#client_certificate[client_certificate]
a|Client certificate (name and UUID)


|server_ca_certificates
|array[link:#server_ca_certificates[server_ca_certificates]]
a|The array of certificates that are common for all the keyservers per
SVM.


|servers
|array[link:#key_server_readcreate[key_server_readcreate]]
a|The set of external key servers.


|===
```

```
[#onboard]
[.api-collapsible-fifth-title]
onboard

Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Is the onboard key manager enabled?


|existing_passphrase
|string
a|The cluster-wide passphrase. This is not audited.


|key_backup
|string
a|Backup of the onboard key manager's key hierarchy. It is required to
save this backup after configuring the onboard key manager to help in the
recovery of the cluster in case of catastrophic failures.


|===


[#status]
[.api-collapsible-fifth-title]
status

Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns 0 if the setup is
complete. This is an advanced property; there is an added computational
cost to retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|message
|string
a|Current state of the key manager indicating any additional steps to
perform to finish the setup. This is an advanced property; there is an
added computational cost to retrieving its value. The property is not
populated for either a collection GET or an instance GET unless it is
explicitly requested using the `fields` query parameter or GET for all
advanced properties is enabled.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume_encryption]
[.api-collapsible-fifth-title]
volume_encryption

Indicates whether volume encryption is supported in the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns a 0 if volume
encryption is supported in all nodes of the cluster.


|message
|string
a|Reason for not supporting volume encryption.


|supported
|boolean
a|Set to true when volume encryption support is available on all nodes of
the cluster.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[IDdfbe73e42bc767ec3c3922e46aa6e77a]]
= Update key managers

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/key-managers/{uuid}`#

*Introduced In:* 9.6

Updates a key manager.
*Note:* When validate_certificates is set, the API returns a 202 status
code for successful creation of the job, otherwise an error is returned.

== Required properties (when patching the Onboard Key Manager)

* `onboard.existing_passphrase` - Cluster-wide passphrase. Required only
when synchronizing the passphrase of the Onboard Key Manager.
* `synchronize` - Synchronizes missing Onboard Key Manager keys on any
node in the cluster. Required only when synchronizing the Onboard Key
Manager keys in a local cluster.

== Required properties (when patching an external key manager)

* `external.client_certificate` or `external.server_ca_certificates` -
Client certificate or Server CA certificate. Required when modifying an
external key manager.

== Optional parameters (when patching an external key manager):

* `validate_certificates` - Validate certificates by testing the
connection to the KMIP server using them. This flag is set to "false" by
default.

== Related ONTAP commands

* `security key-manager external modify`
* `security key-manager onboard sync`
* `security key-manager onboard update-passphrase`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Key manager UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|is_default_data_at_rest_encryption_disabled
|boolean
a|Indicates whether default data-at-rest encryption is disabled in the
cluster. This field is deprecated in ONTAP 9.8 and later. Use the
"software_data_encryption.disabled_by_default" of /api/security endpoint.

* Default value: 1
* Introduced in: 9.7
* x-ntap-readModify: true
* x-nullable: true


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


|uuid
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "external": {
    "client_certificate": {
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
```

```
    "server_ca_certificates": [
      {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  },
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII
characters.",
    "key_backup": "'--------------------------BEGIN
BACKUP--------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAAAxBFWWAAAAACEAAAAAAAAA
QAAAAAAAAABzDyyVAAAAALI5Jsjvy6gUxnT78KoDKXHYb6sSeraM00quOULY6BeV
n6dMFxuErCD1lbERaOQZSuaYy1p8oQHtTEfGMLZM4TYiAAAAAAAAACgAAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAAA
IgAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/LRzU
QRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAAADV1Vd/AAAAAMFM9Q229Bhp
mDaTSdqku5DCd8wG+fOZSr4bx4JT5WHvV/r5gJnXDQQAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJsb2IA
AQAAAAMAAAAYAQAAAAAALgePkcAAAAAIgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAA
AAAAAAAAAAACAAAAAAAJAGr3tJA/LRzUQRHwv+1aWvAAAAAAAAAAACIAAAAAAAAA
KAAAAAAAAACIlCHZAAAAAAAAAAAAAAAAAgAAAAAAAQCafcabsxRXMM7gxhLRrzxh
AAAAAAAAAAAkAAAAAAAAAIAAAAAAAAAA2JjQBQAAAACt4IqXcNpVggahl0axLsN4
yQjnNVKWY7mANB29O42hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAAIAAAAAAAkA
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAAIgAAAAAAAAoAAAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAAACQAAAAAAAA
gAAAAAAAAACnMmUtAAAAAGVk8AtPzENFgsGdsFvnmucmYrlQCsFew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
---------------------------END BACKUP---------------------------'",
    "passphrase": "The cluster password of length 32-256 ASCII
characters."
  },
  "uuid": "string"
}
```

```
====

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536139
| The existing passphrase value provided does not match the configured
passphrase.

| 65536150
| The new passphrase is same as old passphrase.

| 65536404
| The passphrase does not match the accepted length.

| 65536406
| The change of passphrase failed.

| 65536407
| The passphrase update failed on some nodes.

| 65536802
| The passphrase does not match the accepted length in common criteria
mode.

| 65536821
| The certificate is not installed.

| 65536828
| External key management is not enabled for the SVM.
```

```
| 65536850
| New client certificate public or private keys are different from the
existing client certificate.

| 65536852
| Failed to query supported KMIP protocol versions.

| 65536917
| Updating an onboard passphrase requires both new and existing cluster
passphrase.

| 65537242
| The Onboard Key Manager existing_passphrase must be provided when
performing a PATCH/synchronize operation.

| 65537243
| The Onboard Key Manager passphrase must not be provided when performing
a PATCH/synchronize operation.

| 65538120
| The key manager policy is not supported on the admin SVM.

| 65539586
| Cannot modify an inactive key manager configuration.

| 66060338
| Failed to establish secure connection for a key management server due to
incorrect server_ca certificates.

| 66060339
| Failed to establish secure connection for a key management server due to
incorrect client certificates.

| 66060340
| Failed to establish secure connection for a key management server due to
Cryptsoft error.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.
```

```
|===


[#client_certificate]
[.api-collapsible-fifth-title]
client_certificate

Client certificate (name and UUID)


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name


|uuid
|string
a|Certificate UUID


|===


[#server_ca_certificates]
[.api-collapsible-fifth-title]
server_ca_certificates

Security certificate object reference


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID


|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
```

```
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#key_server_readcreate]
[.api-collapsible-fifth-title]
key_server_readcreate
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|secondary_key_servers
|string
a|A comma delimited string of the secondary key servers associated with
the primary key server.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|Username credentials for connecting with the key server.


|===


[#external]
[.api-collapsible-fifth-title]
external

Configures external key management


[cols=3*,options=header]
|===
|Name
|Type
|Description

|client_certificate
|link:#client_certificate[client_certificate]
a|Client certificate (name and UUID)


|server_ca_certificates
|array[link:#server_ca_certificates[server_ca_certificates]]
a|The array of certificates that are common for all the keyservers per
SVM.
```

```
|===


[#onboard]
[.api-collapsible-fifth-title]
onboard

Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
that you can use it if you need to perform a manual recovery operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Is the onboard key manager enabled?


|existing_passphrase
|string
a|The cluster-wide passphrase. This is not audited.


|key_backup
|string
a|Backup of the onboard key manager's key hierarchy. It is required to
save this backup after configuring the onboard key manager to help in the
recovery of the cluster in case of catastrophic failures.


|passphrase
|string
a|The cluster-wide passphrase. This is not audited.


|synchronize
|boolean
a|Synchronizes missing onboard keys on any node in the cluster. If a node
is added to a cluster that has onboard key management configured, the
synchronize operation needs to be performed in a PATCH operation. In a
```

```
MetroCluster configuration, if onboard key management is enabled on one
site, then the synchronize operation needs to be run as a POST operation
on the remote site providing the same passphrase.


|===


[#status]
[.api-collapsible-fifth-title]
status

Optional status information on the current state of the key manager
indicating if it is fully setup or requires more action.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns 0 if the setup is
complete. This is an advanced property; there is an added computational
cost to retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|message
|string
a|Current state of the key manager indicating any additional steps to
perform to finish the setup. This is an advanced property; there is an
added computational cost to retrieving its value. The property is not
populated for either a collection GET or an instance GET unless it is
explicitly requested using the `fields` query parameter or GET for all
advanced properties is enabled.


|===


[#svm]
[.api-collapsible-fifth-title]
```

```
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume_encryption]
[.api-collapsible-fifth-title]
volume_encryption

Indicates whether volume encryption is supported in the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Code corresponding to the status message. Returns a 0 if volume
encryption is supported in all nodes of the cluster.


|message
|string
a|Reason for not supporting volume encryption.
```

```
|supported
|boolean
a|Set to true when volume encryption support is available on all nodes of
the cluster.


|===


[#security_key_manager]
[.api-collapsible-fifth-title]
security_key_manager

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether the configuration is enabled.


|external
|link:#external[external]
a|Configures external key management


|is_default_data_at_rest_encryption_disabled
|boolean
a|Indicates whether default data-at-rest encryption is disabled in the
cluster. This field is deprecated in ONTAP 9.8 and later. Use the
"software_data_encryption.disabled_by_default" of /api/security endpoint.

* Default value: 1
* Introduced in: 9.7
* x-ntap-readModify: true
* x-nullable: true


|onboard
|link:#onboard[onboard]
a|Configures onboard key management. After configuring onboard key
management, save the encrypted configuration data in a safe location so
```

```
that you can use it if you need to perform a manual recovery operation.


|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage authentication keys

:leveloffset: +1


[[IDe9da61c9e0355dd8c0686ec2cbbe54ab]]
= Manage authentication keys



:doctype: book

== Overview

Used to create and list authentication keys (NSE-AK).

== Example

=== Creates an authentication key.
```

```
NOTE: an external key  manager must be configured on the admin SVM, as the
authentication key will be associated with the admin SVM.
The following example shows how to create an authentication key.

----

# The API:
POST /api/security/key-managers/{security_key_manager.uuid}/auth-keys

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers/5fb1701a-d922-
11e8-bfe8-005056bb017d/auth-keys?return_timeout=15&return_records=true' -H
'accept: application/hal+json'

# The response:
{
"num_records": 1,
"records": [
  {
    "key_id":
"00000000000000000200000000000100531d8cdc38437c2627b6b1726dd2675c000000000
0000000",
    "key_tag": "vsim1"
  }
]
}
----

'''

=== Retrieving a list of all of the authentication keys associated with
the admin SVM.

The following example shows how to retrieve a list of all of the
authentication keys associated with the admin SVM.

----

# The API:
GET /api/security/key-managers/{security_key_manager.uuid}/auth-keys

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/5fb1701a-d922-
11e8-bfe8-005056bb017d/auth-keys?fields=*' -H 'accept:
application/hal+json'

# The response:
```

```
{
"num_records": 2,
"records": [
  {
    "security_key_manager": {
      "uuid": "d36a654d-14b4-11ed-b82e-005056bb8f59"
    },
    "key_id":
"00000000000000000200000000000100052ab79fc51a430dbb16f1c0d2054cfe000000000
0000000",
    "key_tag": "vsim1"
  },
  {
    "security_key_manager": {
      "uuid": "d36a654d-14b4-11ed-b82e-005056bb8f59"
    },
    "key_id":
"00000000000000000200000000000001003f32ce2dc55d2764c07da74e722c179b000000000
0000000",
    "key_tag": "vsim1"
  }
]
}
----

'''

=== Retrieving a specific authentication key associated with the admin
SVM.

The following example shows how to a specific authentication key
associated with the admin SVM and return the key-tag.
```

**The API:**

GET /api/security/key-managers/{security_key_manager.uuid}/auth-keys/{key-id}

**The call:**

curl -X GET 'https://<mgmt-ip>/api/security/key-managers/5fb1701a-d922-11e8-bfe8-005056bb017d/auth-keys/00000000000000000200000000000010041a2dda969b0d179db8f1c78d629d0f10000000000000000?fields=key_tag' -H 'accept: application/hal+json'</mgmt-ip>

**The response:**

{
"security_key_manager": {
"uuid": "d36a654d-14b4-11ed-b82e-005056bb8f59"

},
"key_id": "00000000000000000200000000000010041a2dda969b0d179db8f1c78d629d0f10000000000000000",
"key_tag": "vsim1"
}

### Retrieve authentication keys associated with the admin SVM

GET `/security/key-managers/{security_key_manager.uuid}/auth-keys`

**Introduced In:** 9.12

Retrieves a list of all authentication keys associated with the admin SVM.

#### Related ONTAP commands

- `security key-manager key query`

#### Required properties

- `security_key_manager.uuid` - UUID of the external key manager.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| key_tag | string | query | False | Filter by key_tag |
| key_id | string | query | False | Filter by key_id |
| security_key_manager.uuid | string | path | True | External key manager UUID |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[key_manager_auth_key] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "key_id":
"00000000000000000200000000000001003aa8ce6a4fea3e466620134bea9510a1000000
0000000000",
      "key_tag": "Authentication-Key-01",
      "security_key_manager": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      }
    }
  ]
}
```


#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65536828 | External key management is not enabled for the SVM. |
| 65536856 | No key servers found for the SVM. |

| Error Code | Description |
| --- | --- |
| 65536896 | External key management is not configured on the partner site. |
| 65538800 | External KMIP DKMIP keymanager not configured on administrative Vserver. |
| 65538801 | Internal error while accessing keymanager database. |
| 65538802 | The UUID provided is not associated with the administrator SVM key manager. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

security_key_manager

Security key manager object reference.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | Key manager UUID |

key_manager_auth_key

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| key_id | string | Key identifier. |
| key_tag | string | Optional parameter to define key-tag for the authentication key, length 0-32 characters. |
| security_key_manager | security_key_manager | Security key manager object reference. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an authentication key

POST `/security/key-managers/{security_key_manager.uuid}/auth-keys`

**Introduced In:** 9.12

Creates an authentication key.

#### Related ONTAP commands

- `security key-manager key create`

#### Required properties

- `security_key_manager.uuid` - UUID of the external key manager.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_key_manager.uuid | string | path | True | External key manager UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| key_id | string | Key identifier. |
| key_tag | string | Optional parameter to define key-tag for the authentication key, length 0-32 characters. |
| passphrase | string | Authentication passphrase, length 20-32 characters. May contain the '=' character. |
| security_key_manager | security_key_manager | Security key manager object reference. |

#### Example request

```
{
  "key_id":
"0000000000000000020000000000001003aa8ce6a4fea3e466620134bea9510a1000000
0000000000",
  "key_tag": "Authentication-Key-01",
  "passphrase": "AuthenticationKey_01",
  "security_key_manager": {
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[key_manager_auth_key] | |

#### Example response

```
{
  "records": [
    {
      "key_id":
"0000000000000000020000000000001003aa8ce6a4fea3e466620134bea9510a1000000
0000000000",
      "key_tag": "Authentication-Key-01",
      "passphrase": "AuthenticationKey_01",
      "security_key_manager": {
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      }
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65536048 | The maximum number of authentication keys are already configured. |
| 65536053 | Invalid passphrase length; passphrase must be 20-32 ASCII-range characters. |
| 65536056 | The key tag value provided contains invalid characters. |

| Error Code | Description |
| --- | --- |
| 65536056 | The key-tag option cannot contain any spaces, tabs or new lines. |
| 65536076 | Failed to push authentication key to any registered key servers. |
| 65536160 | Unable to determine the current number of configured authentication keys. |
| 65536264 | Failed to create authentication key. |
| 65536265 | Failed to create a key-id for the authentication key. |
| 65536828 | External key management is not enabled for the SVM. |
| 65536856 | No key servers found for the SVM. |
| 65536872 | Error cleaning up key database after key creation error. |
| 65536896 | External key management is not configured on the partner site. |
| 65536987 | One or more key servers are unavailable. |
| 65538800 | External KMIP DKMIP keymanager not configured on administrative Vserver. |
| 65538801 | Internal error while accessing keymanager database. |
| 65538802 | The UUID provided is not associated with the administrator SVM key manager. |
| 66060289 | Failed to store authentication key on key server. |
| 66060304 | Invalid key length. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

_links

security_key_manager

Security key manager object reference.

| Name | Type | Description |
|---|---|---|
| uuid | string | Key manager UUID |

key_manager_auth_key

| Name | Type | Description |
|---|---|---|
| key_id | string | Key identifier. |
| key_tag | string | Optional parameter to define key-tag for the authentication key, length 0-32 characters. |
| passphrase | string | Authentication passphrase, length 20-32 characters. May contain the '=' character. |
| security_key_manager | security_key_manager | Security key manager object reference. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete an authentication key

DELETE `/security/key-managers/{security_key_manager.uuid}/auth-keys/{key_id}`

**Introduced In:** 9.12

Deletes an authentication key.

#### Related ONTAP commands

- `security key-manager key delete`

#### Required properties

- `security_key_manager.uuid` - UUID of the external key manager.
- `key_id` - Key ID of the authentication key to be deleted.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| security_key_manager.uuid | string | path | True | External key manager UUID |
| key_id | string | path | True | Key ID of the authentication key to be deleted. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65536828 | External key management is not enabled for the SVM. |
| 65536859 | Authentication key-Id provided for deletion is in use. |
| 65536860 | Key-id provided for deletion is not an authentication key. |
| 65538800 | External KMIP DKMIP keymanager not configured on administrative Vserver. |
| 65538801 | Internal error while accessing keymanager database. |
| 65538802 | The UUID provided is not associated with the administrator SVM key manager. |
| 66060296 | Failed to delete key from an external key server. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### List all authentication keys associated with the admin SVM

GET `/security/key-managers/{security_key_manager.uuid}/auth-keys/{key_id}`

**Introduced In:** 9.12

Retrieves the authentication key identified by the 'key_id' and associated with the admin SVM.

#### Related ONTAP commands

- `security key-manager key query`

#### Required properties

- `security_key_manager.uuid` - UUID of the external key manager.
- `key_id` - Key ID of the authentication key to be retrieved.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| key_tag | string | query | False | Filter by key_tag |
| security_key_manager.uuid | string | path | True | External key manager UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| key_id | string | path | True | Key ID of the authentication key to be retrieved. |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
|---|---|---|
| key_id | string | Key identifier. |
| key_tag | string | Optional parameter to define key-tag for the authentication key, length 0-32 characters. |
| security_key_manager | security_key_manager | Security key manager object reference. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "key_id":
"0000000000000000020000000000001003aa8ce6a4fea3e466620134bea9510a1000000
0000000000",
  "key_tag": "Authentication-Key-01",
  "security_key_manager": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  }
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 65536828 | External key management is not enabled for the SVM. |
| 65536856 | No key servers found for the SVM. |

| Error Code | Description |
| --- | --- |
| 65536896 | External key management is not configured on the partner site. |
| 65538800 | External KMIP DKMIP keymanager not configured on administrative Vserver. |
| 65538801 | Internal error while accessing keymanager database. |
| 65538802 | The UUID provided is not associated with the administrator SVM key manager. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

security_key_manager

Security key manager object reference.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | Key manager UUID |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# View security key managers

## View security key managers

### Overview

Retrieves the key manager keys on the give node. The following operations are supported:

- Get

### Examples

## Retrieving key manager key-id information for a node

The following example shows how to retrieve key-ids present on a node for a key manager.

```
# The API:
GET /api/security/key-
managers/{security_key_manager.uuid}/keys/{node.uuid}/key-ids

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/f4f98a48-8a5c-
c548-cd03-c6335f5803a8/keys/00000000-0000-0000-0000-000000000000/key-
ids/0000000000000000020000000000005009ad4da8fea2cafe2bed803078b780ebe000000
0000000c01' -H 'accept: application/hal+json'

# The response:
{
  "security_key_manager": {
    "uuid": "f4f98a48-8a5c-c548-cd03-c6335f5803a8"
  },
  "node": {
    "uuid": "00000000-0000-0000-0000-000000000000"
  },
  "key_id":
"0000000000000000020000000000005009ad4da8fea2cafe2bed803078b780ebe000000000
0000c01",
  "svm": {
    "name": "cluster-1"
  },
  "key_tag": "vsim1",
  "key_type": "vek",
  "restored": false,
  "key_store": "onboard",
  "key_user": "datavs",
  "key_manager": "onboard",
  "key_store_type": "okm",
  "encryption_algorithm": "XTS-AES-256",
  "_links": {
    "self": {
      "href": "/api/security/key-managers/f4f98a48-8a5c-c548-cd03-
c6335f5803a8/keys/00000000-0000-0000-0000-000000000000/key-
ids/0000000000000000020000000000005009ad4da8fea2cafe2bed803078b780ebe000000
0000000c01"
    }
  }
}
```

#### Retrieving key manager key-id information of a specific key-type for a node

The following example shows how to retrieve key-ids of a specific key-type present on a node for a key manager.

```
= The API:

GET /api/security/key-
manager/{security_key_manager.uuid}/keys/{node.uuid}/key-ids

= The call:

curl -X GET "https://+++<mgmt-ip>+++/api/security/key-managers/7c179931-
044b-11ed-b7cd-005056bbc535/keys/44dac31e-0449-11ed-b7cd-005056bbc535/key-
ids?key_type=nse_ak&return_records=true&return_timeout=15" -H "accept:
application/json"+++</mgmt-ip>+++

= The response:

{
  "records": [
    {
      "key_server": "10.225.89.34:5696",
      "key_id":
"00000000000000000200000000000001003d5c5f8c497e8e36aa80566e08749a3d000000000
0000000",
      "key_type": "nse_ak"
    },
    {
      "key_server": "10.225.89.34:5696",
      "key_id":
"00000000000000000200000000000100c2dce9a3a15aeb8480db8d49c17d056c000000000
0000000",
      "key_type": "nse_ak"
    }
  ],
  "num_records": 2
}





[[ID70bfe43ce561fb4458162033a07dcb00]]
= Retrieve key manager configurations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/security/key-
managers/{security_key_manager.uuid}/keys/{node.uuid}/key-ids`#

*Introduced In:* 9.11

Retrieves key manager configurations.

== Required properties

* `security_key_manager.uuid` - Key manager UUID.
* `node.uuid` - Node UUID.

== Related ONTAP commands

* `security key-manager key query`
* `security key-manager key query -node <node>`
* `security key-manager key query -node <node> -key-manager <key_manager>`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|key_server
|string
|query
|False
a|Filter by key_server

* Introduced in: 9.12


|encryption_algorithm
|string
|query
|False
a|Filter by encryption_algorithm

* Introduced in: 9.12
```

```
|node.name
|string
|query
|False
a|Filter by node.name

* Introduced in: 9.12


|policy
|string
|query
|False
a|Filter by policy

* Introduced in: 9.12


|key_tag
|string
|query
|False
a|Filter by key_tag

* Introduced in: 9.12


|key_manager
|string
|query
|False
a|Filter by key_manager

* Introduced in: 9.12


|svm.name
|string
|query
|False
a|Filter by svm.name

* Introduced in: 9.14


|svm.uuid
```

```
|string
|query
|False
a|Filter by svm.uuid

* Introduced in: 9.14


|crn
|string
|query
|False
a|Filter by crn

* Introduced in: 9.12


|key_store_type
|string
|query
|False
a|Filter by key_store_type

* Introduced in: 9.12


|key_type
|string
|query
|False
a|Filter by key_type

* Introduced in: 9.12


|key_user
|string
|query
|False
a|Filter by key_user

* Introduced in: 9.12


|key_store
|string
|query
```

```
|False
a|Filter by key_store

* Introduced in: 9.12


|restored
|boolean
|query
|False
a|Filter by restored

* Introduced in: 9.12


|scope
|string
|query
|False
a|Filter by scope

* Introduced in: 9.14


|key_id
|string
|query
|False
a|Filter by key_id

* Introduced in: 9.12


|security_key_manager.uuid
|string
|path
|True
a|Key manager UUID.


|node.uuid
|string
|path
|True
a|Node UUID.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#key_manager_keys[key_manager_keys]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
      },
      "crn": "CRN=v1:bluemix:public:containers-kubernetes",
      "encryption_algorithm": "XTS-AES-256",
      "key_id":
"0000000000000000020000000000001008963c9213194c59555c1bec8db3603c800000000"
,
      "key_manager": "keyserver1.local:5696",
      "key_server": "keyserver1.com:5698",
      "key_store": "string",
      "key_store_type": "string",
      "key_tag": "key#",
      "key_type": "string",
      "key_user": "vs1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "policy": "IBM_Key_Lore",
      "restored": 1,
      "scope": "string",
      "security_key_manager": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====
```

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65537125
| The provided key manager UUID is not the UUID of a keymanager.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|
```

```
|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#security_key_manager]
[.api-collapsible-fifth-title]
security_key_manager

Security key manager object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|Key manager UUID


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#key_manager_keys]
[.api-collapsible-fifth-title]
key_manager_keys

Displays the keys stored in a key manager.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|crn
|string
a|Cloud resource name.
```

```
|encryption_algorithm
|string
a|Encryption algorithm for the key


|key_id
|string
a|Key identifier.


|key_manager
|string
a|Key manager key server managing the key. Indicates the external key
server when external key manager is configured.


|key_server
|string
a|External key server for key management.


|key_store
|string
a|Security key manager configured for the given key manager UUID. Key
manager keystore value can be onboard or external.


|key_store_type
|string
a|Security key manager keystore type. Keystore type can be onboard,
external, or supported cloud key manager.


|key_tag
|string
a|Additional information associated with the key.


|key_type
|string
a|Encryption Key type.


|key_user
|string
a|SVM associated with the key.
```

```
|node
|link:#node[node]
a|

|policy
|string
a|Key store policy.


|restored
|boolean
a|Indicates whether the key is present locally on the node.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|security_key_manager
|link:#security_key_manager[security_key_manager]
a|Security key manager object reference.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID8158c84eb160e141bf2dac78ffa77d4c]]
= Retrieve key management key information for specified keys

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-
managers/{security_key_manager.uuid}/keys/{node.uuid}/key-ids/{key_id}`#

*Introduced In:* 9.11

Retrieves the key management keys information for the specified key_id.

== Related ONTAP commands

* `security key-manager key query -key-id <key_id>`
* `security key-manager key query -key-id <key_id> -node <node>`
* `security key-manager key query -key-id <key_id> -node <node> -key
-manager <key_manager>`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|security_key_manager.uuid
|string
|path
|True
a|Key manager UUID.


|node.uuid
|string
|path
|True
a|Node UUID.


|key_id
|string
```

```
|path
|True
a|Key identifier.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|crn
|string
a|Cloud resource name.


|encryption_algorithm
|string
a|Encryption algorithm for the key


|key_id
|string
a|Key identifier.


|key_manager
|string
a|Key manager key server managing the key. Indicates the external key
server when external key manager is configured.
```

```
|key_server
|string
a|External key server for key management.


|key_store
|string
a|Security key manager configured for the given key manager UUID. Key
manager keystore value can be onboard or external.


|key_store_type
|string
a|Security key manager keystore type. Keystore type can be onboard,
external, or supported cloud key manager.


|key_tag
|string
a|Additional information associated with the key.


|key_type
|string
a|Encryption Key type.


|key_user
|string
a|SVM associated with the key.


|node
|link:#node[node]
a|

|policy
|string
a|Key store policy.


|restored
|boolean
a|Indicates whether the key is present locally on the node.
```

```
|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|security_key_manager
|link:#security_key_manager[security_key_manager]
a|Security key manager object reference.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "crn": "CRN=v1:bluemix:public:containers-kubernetes",
  "encryption_algorithm": "XTS-AES-256",
  "key_id":
"0000000000000000020000000000001008963c9213194c59555c1bec8db3603c800000000"
,
  "key_manager": "keyserver1.local:5696",
  "key_server": "keyserver1.com:5698",
  "key_store": "string",
  "key_store_type": "string",
  "key_tag": "key#",
  "key_type": "string",
  "key_user": "vs1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "policy": "IBM_Key_Lore",
  "restored": 1,
  "scope": "string",
  "security_key_manager": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|
```

```
|uuid
|string
a|

|===


[#security_key_manager]
[.api-collapsible-fifth-title]
security_key_manager

Security key manager object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|Key manager UUID


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDc7c107eccc06948b0f7d1f4b19d9e6a4]]
= Retrieve and restore unrestored keys

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/key-managers/{security_key_manager.uuid}/restore`#

*Introduced In:* 9.13

Retrieves and restores any current unrestored keys (associated with the
storage controller) from the specified key management server.

== Required properties

* `security_key_manager.uuid` - UUID of the key management server.
The UUID of the external key manager can be retrieved using [`GET
/api/security/key-managers`].
```

```
== Related ONTAP commands

* `security key-manager external restore`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|security_key_manager.uuid
|string
|path
|True
a|Key manager UUID.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.
```

```
* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier.


|key_server
|string
a|External key server for key management.


|key_tag
|string
a|Additional information associated with the key.


|node
|link:#node[node]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "key_id":
"0000000000000000020000000000001003aa8ce6a4fea3e466620134bea9510a1000000000
0000000",
  "key_server": "keyserver1.com:5698",
  "key_tag": "Authentication-Key-01",
  "node": {
```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 65536083
| Internal error. Failed to restore the authentication key.

| 65536843
| The key management server is not configured for the SVM.

| 65536855
| Internal error. Failed to restore the volume encryption key.

| 65538500
| This command only restores keys from primary key servers.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|

|===


[#security_key_manager_restore_keys]
[.api-collapsible-fifth-title]
security_key_manager_restore_keys

[cols=3*,options=header]
|===
|Name
|Type
|Description

|key_id
|string
a|Key identifier.


|key_server
|string
a|External key server for key management.


|key_tag
|string
a|Additional information associated with the key.


|node
|link:#node[node]
a|

|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


[[ID31a1f910738357b98c8924dd017fed1c]]
= Migrate SVM keys between security key managers

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/key-managers/{source.uuid}/migrate`#

*Introduced In:* 9.7

Migrates the keys belonging to an SVM between the cluster's key manager
```

```
and the SVM's key manager. This operation can run for several minutes.

== Required properties

* `source.uuid` - UUID of the source key manager.
* `uuid` - UUID of the destination key manager.
The UUID of onboard and external KMIP key manager can be fetched using
[`GET /api/security/key-managers`].
The UUID of Azure Key Vault key manager can be fetched using [`GET
/api/security/azure-key-vaults`].
The UUID of Google Cloud key manager can be fetched using [`GET
/api/security/gcp-kms`].
The UUID of Amazon Web Services key manager can be fetched using [`GET
/api/security/aws-kms`].

== Related ONTAP commands

* `security key-manager key migrate`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|source.uuid
|string
|path
|True
a|Migration source key manager UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
```

```
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|Key manager UUID


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536886
| The specified migration option is not supported in this release.

| 65536959
| The source-uuid and UUID must be different values.
```

```
| 65536968
| Check that all nodes of the cluster are healthy and retry the operation.

| 65537117
| The migrate operation cannot be started because a UUID cannot be
converted to an SVM name.

| 65537117
| Cannot start migration because a key manager referenced by a provided
UUID does not exist.

| 65537551
| Top-level internal key protection key (KEK) is unavailable on one or
more nodes.

| 65537552
| Embedded KMIP server status is not available.

| 65537564
| Check that the Azure Key Vault Service is healthy and retry the
operation.

| 65537611
| Key migration cannot be performed on the SVM while the enabled keystore
configuration is being switched. If a previous attempt to switch the
keystore configuration failed, or was interrupted, the system will
continue to prevent key migration for the SVM. Use the REST API PATCH
method "/api/security/key-stores/\{uuid}" to re-run and complete the
operation.

| 65537720
| Failed to configure the Google Cloud Key Management Service for an SVM
because a key manager is already configured.

| 65537736
| Check that the Google Cloud Key Management Service is healthy and retry
the operation.

| 65538107
| Key migration to an IBM Key Lore key manager is not supported.

| 65539434
| Cannot perform key migration on the SVM while the keystore is being
initialized. Wait until the keystore is in the active state, and rerun the
key migration operation.
```

```
| 65539435
| Cannot perform key migration on the SVM while the keystore is being
disabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#info]
[.api-collapsible-fifth-title]
info

Migration destination key manager UUID


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|uuid
|string
a|Key manager UUID


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID609f295102ab5235ffa56104faa14335]]
= List key servers configured in an external key manager
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-managers/{uuid}/key-servers`#

*Introduced In:* 9.6

Retrieves the list of key servers configured in an external key manager.

== Expensive properties

There is an added computational cost to retrieving these properties. They
are excluded from default GET results and must be explicitly requested
using the `fields` query parameter. For more details, see
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields]. Additionally, these fields
are unavailable for inactive configurations as they are only relevant to
active configurations.

* `connectivity.cluster_availability`
* `connectivity.node_states.node.name`
* `connectivity.node_states.node.uuid`
* `connectivity.node_states.state`

== Examples

* To retrieve basic information about a key server:
+
----
GET /security/key-managers/{uuid}/key-servers
----

* To retrieve specific fields, including expensive properties:
+
----
GET /security/key-managers/{uuid}/key-
servers?fields=connectivity.cluster_availability,connectivity.node_states.
node.name
----

== Related ONTAP commands

* `security key-manager external show`
* `security key-manager external show-status`


== Parameters

[cols=5*,options=header]
```

```
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|External key manager UUID


|secondary_key_servers
|string
|query
|False
a|Filter by secondary_key_servers

* Introduced in: 9.8


|username
|string
|query
|False
a|Filter by username


|server
|string
|query
|False
a|Filter by server


|timeout
|integer
|query
|False
a|Filter by timeout

* Max value: 60
* Min value: -1
```

```
|create_remove_timeout
|integer
|query
|False
a|Filter by create_remove_timeout

* Introduced in: 9.14
* Max value: 60
* Min value: -1


|connectivity.node_states.node.name
|string
|query
|False
a|Filter by connectivity.node_states.node.name

* Introduced in: 9.13


|connectivity.node_states.node.uuid
|string
|query
|False
a|Filter by connectivity.node_states.node.uuid

* Introduced in: 9.13


|connectivity.node_states.state
|string
|query
|False
a|Filter by connectivity.node_states.state

* Introduced in: 9.13


|connectivity.cluster_availability
|boolean
|query
|False
a|Filter by connectivity.cluster_availability

* Introduced in: 9.7
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#key_server[key_server]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
```

```
        }
      },
      "connectivity": {
        "node_states": [
          {
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "state": "not_responding"
          }
        ]
      },
      "create_remove_timeout": 60,
      "password": "password",
      "secondary_key_servers": [
        "secondary1.com",
        "10.1.2.3"
      ],
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
```

```
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|KMIP username credentials for connecting with the key server.


|===


[#key_server]
[.api-collapsible-fifth-title]
key_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
```

```
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|create_remove_timeout
|integer
a|The key server timeout for create and remove operations.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|secondary_key_servers
|array[string]
a|A list of the secondary key servers associated with the primary key
server.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|username
|string
a|KMIP username credentials for connecting with the key server.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID22d1444fc6f8cfab7bcdecfa01c57175]]
= Add primary key servers to an external key manager

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/key-managers/{uuid}/key-servers`#

*Introduced In:* 9.6

Adds primary key servers to a configured external key manager.

== Required properties

* `uuid` - UUID of the external key manager.
* `server` - Primary Key server name.

== Related ONTAP commands

* `security key-manager external add-servers`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|External key manager UUID


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|records
|array[link:#records[records]]
a|An array of key servers specified to add multiple key servers to a key
manager in a single API call. Valid in POST only and not valid if `server`
is provided.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.
```

```
|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "password": "password",
  "records": [
    {
      "password": "password",
      "server": "bulkkeyserver.com:5698"
    }
  ],
  "server": "keyserver1.com:5698"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#key_server[key_server]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "num_records": 1,
  "records": [
    {
      "password": "password",
      "records": [
        {
          "password": "password",
          "server": "bulkkeyserver.com:5698"
        }
      ],
      "server": "keyserver1.com:5698"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536038
| A maximum of 4 active primary key servers are allowed.
```

```
| 65536042
| Cannot add key server because it is already a secondary key server.

| 65536600
| Cannot add a key server while a node is out quorum.

| 65536821
| The certificate is not installed.

| 65536824
| Multitenant key management is not supported in MetroCluster
configurations.

| 65536828
| External key management is not enabled for the SVM.

| 65536834
| Failed to get existing key-server details for the SVM.

| 65536852
| Failed to query supported KMIP protocol versions.

| 65536870
| Key management servers are already configured.

| 65536870
| The key management servers already exist.

| 65536871
| Duplicate key management servers exist.

| 65536921
| Unable to execute the command on the KMIP server.

| 66060338
| Unable to establish secure connection to KMIP server due to incorrect
server_ca certificates.

| 66060339
| Unable to establish secure connection to KMIP server due to incorrect
client certificates.

| 66060340
| Unable to establish secure connection to KMIP server due to Cryptsoft
error.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
```

```
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|===


[#key_server]
[.api-collapsible-fifth-title]
key_server

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|records
|array[link:#records[records]]
a|An array of key servers specified to add multiple key servers to a key
manager in a single API call. Valid in POST only and not valid if `server`
is provided.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0236b1e802a30e9362ce1476054fc4f4]]
= Delete a primary key server

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/key-managers/{uuid}/key-servers/{server}`#
```

```
*Introduced In:* 9.6

Deletes a primary key server.

== Optional parameters:

* `force` - Bypass Out of Quorum checks when deleting a primary key
server. This flag is set to "false" by default.

== Related ONTAP commands

* `security key-manager external remove-servers`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|External key manager UUID


|server
|string
|path
|True
a|Primary key server configured in the external key manager.


|force
|boolean
|query
|False
a|Set the force flag to "true" to bypass out of quorum checks when
removing a primary key server.

* Introduced in: 9.11
```

```
* Default value:

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536600
| Cannot remove a key server while a node is out of quorum.

| 65536700
| The key server contains keys that are currently in use and not available
from any other configured key server in the SVM.

| 65536843
| The key management server is not configured for the SVM.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID152be7951c1634f681570f0717b3b451]]
= Retrieve key servers configured in an external key manager

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-managers/{uuid}/key-servers/{server}`#
```

```
*Introduced In:* 9.6

Retrieves key servers configured in an external key manager.

== Expensive properties

There is an added computational cost to retrieving these properties. They
are excluded from default GET results and must be explicitly requested
using the `fields` query parameter. For more details, see
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields]. Additionally, these fields
are unavailable for inactive configurations as they are only relevant to
active configurations.

* `connectivity.cluster_availability`
* `connectivity.node_states.node.name`
* `connectivity.node_states.node.uuid`
* `connectivity.node_states.state`

== Examples

* To retrieve basic information about a key server:
+
----
GET /security/key-managers/{uuid}/key-servers/{server}
----

* To retrieve specific fields, including expensive properties:
+
----
GET /security/key-managers/{uuid}/key-
servers/{server}?fields=connectivity.cluster_availability,connectivity.nod
e_states.node.name
----

== Related ONTAP commands

* `security key-manager external show`
* `security key-manager external show-status`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|External key manager UUID


|server
|string
|path
|True
a|Primary Key server configured in the key manager.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
the cluster.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|create_remove_timeout
|integer
a|The key server timeout for create and remove operations.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|secondary_key_servers
|array[string]
a|A list of the secondary key servers associated with the primary key
server.


|server
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|username
|string
a|KMIP username credentials for connecting with the key server.


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "connectivity": {
    "node_states": [
      {
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "state": "not_responding"
      }
    ]
  },
  "create_remove_timeout": 60,
  "password": "password",
  "secondary_key_servers": [
    "secondary1.com",
    "10.1.2.3"
  ],
  "server": "keyserver1.com:5698",
  "timeout": 60,
  "username": "username"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|connectivity
|link:#connectivity[connectivity]
a|This property contains the key server connectivity state of all nodes in
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|server
```

```
|string
a|External key server for key management. If no port is provided, a
default port of 5696 is used. Not valid in POST if `records` is provided.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.


|username
|string
a|KMIP username credentials for connecting with the key server.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4f0abb61904a376f9ddae81b0bf18f41]]
= Update a primary key server

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/key-managers/{uuid}/key-servers/{server}`#

*Introduced In:* 9.6

Updates a primary key server.

== Related ONTAP commands

* `security key-manager external modify-server`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|External key manager UUID


|server
|string
|path
|True
a|Primary key server configured in the external key manager.

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|create_remove_timeout
|integer
a|The key server timeout for create and remove operations.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|password
|string
a|Password credentials for connecting with the key server. This is not
```

```
audited.


|secondary_key_servers
|array[string]
a|A list of the secondary key servers associated with the primary key
server.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|username
|string
a|KMIP username credentials for connecting with the key server.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "create_remove_timeout": 60,
  "password": "password",
  "secondary_key_servers": [
    "secondary1.com",
    "10.1.2.3"
  ],
  "timeout": 60,
  "username": "username"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 65536600
| Cannot modify a key server while a node is out quorum.

| 65536824
| Multitenant key management is not supported in MetroCluster
configurations.

| 65536828
| External key management is not enabled for the SVM.

| 65536843
| The key management server is not configured for the SVM.

| 65536845
| Missing username.

| 65536846
| Missing password.

| 65536880
| One or more of the following values must be provided "timeout",
"username", "password", "secondary_key_servers", "create_remove_timeout".

| 65536921
| Unable to execute the command on the KMIP server.

| 65537400
| Exceeded maximum number of secondary key servers.

| 65538407
| A secondary key server is a duplicate of the associated primary key
server.

| 65538408
| The list of secondary key servers contains duplicates.
```

```
| 65538413
| A secondary key server address is not formatted correctly.

| 65538502
| A secondary key server is also a primary key server.

| 65538503
| Support for adding secondary key servers requires an ECV of ONTAP 9.11.1
or later.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#key_server_state]
[.api-collapsible-fifth-title]
key_server_state

The connectivity state of the key server for a specific node.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|state
|string
a|Key server connectivity state


|===


[#connectivity]
[.api-collapsible-fifth-title]
connectivity

This property contains the key server connectivity state of all nodes in
```

```
the cluster.
This is an advanced property; there is an added computational cost to
retrieving its value. The property is not populated for either a
collection GET or an instance GET unless it is explicitly requested using
the `fields` query parameter or GET for all advanced properties is
enabled.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster_availability
|boolean
a|Set to true when key server connectivity state is available on all nodes
of the cluster.


|node_states
|array[link:#key_server_state[key_server_state]]
a|An array of key server connectivity states for each node.


|===


[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.
```

```
|username
|string
a|KMIP username credentials for connecting with the key server.


|===


[#key_server]
[.api-collapsible-fifth-title]
key_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|create_remove_timeout
|integer
a|The key server timeout for create and remove operations.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.


|password
|string
a|Password credentials for connecting with the key server. This is not
audited.


|secondary_key_servers
|array[string]
a|A list of the secondary key servers associated with the primary key
server.


|timeout
|integer
a|I/O timeout in seconds for communicating with the key server.
-1 indicates that the server will wait indefinitely for the event to
occur. 0 indicates that the server will not wait and will immediately
timeout if it does not receive a response.
```

```
|username
|string
a|KMIP username credentials for connecting with the key server.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



= Manage keystores

:leveloffset: +1


[[ID0f91bc76b8fcdb45d2d04d5213424817]]
= Manage keystores



== Overview

A keystore describes a key-manager configuration, specifically the type of
key-manager and whether the configuration is currently enabled for the
configured SVM.

== Examples

'''

=== Retrieving information for all configured key managers

The following example shows how to retrieve information about all
configured key managers.
```

```
----

# The API:
GET /api/security/key-stores

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-stores?fields=*' -H
'accept: application/hal+json'

# The response:
{
"records": [
  {
    "type": "okm",
    "configuration": {
      "name": "default",
      "uuid": "ec8711c9-4e9f-11ef-b477-005056bb677e"
    },
    "uuid": "ec8711c9-4e9f-11ef-b477-005056bb677e",
    "enabled": true,
    "location": "onboard",
    "scope": "cluster",
    "state": "active"
  },
  {
    "type": "kmip",
    "configuration": {
      "name": "default",
      "uuid": "d81f43cd-4e9f-11ef-b477-005056bb677e"
    },
    "uuid": "d81f43cd-4e9f-11ef-b477-005056bb677e",
    "enabled": false,
    "location": "external",
    "scope": "cluster"
  },
  {
    "type": "kmip",
    "configuration": {
      "name": "default",
      "uuid": "7da22185-4ea0-11ef-b477-005056bb677e"
    },
    "uuid": "7da22185-4ea0-11ef-b477-005056bb677e",
    "enabled": true,
    "location": "external",
    "scope": "svm",
```

```
    "svm": {
      "name": "vs0",
      "uuid": "3cbe691b-4ea0-11ef-b477-005056bb677e"
    },
    "state": "active"
  }
],
"num_records": 3
}
----

'''

=== Retrieving a specific keystore by its UUID

The following example shows how to retrieve information about a specific
keystore.

----

# The API:
GET /api/security/key-stores/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-stores/33421d82-0a8d-11ec-
ae88-005056bb5955?fields=*' -H 'accept: application/hal+json'

# The response:
{
"type": "okm",
"configuration": {
  "name": "default",
  "uuid": "ec8711c9-4e9f-11ef-b477-005056bb677e"
},
"uuid": "ec8711c9-4e9f-11ef-b477-005056bb677e",
"enabled": true,
"location": "onboard",
"scope": "cluster",
"state": "active"
}
----

'''

=== Enabling a specific keystore configuration

The following example shows how to enable a specific keystore
```

```
configuration.

----

# The API:
PATCH /api/security/key-stores/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-stores/33421d82-0a8d-
11ec-ae88-005056bb5955' -H 'accept: application/hal+json' -d "{
\"enabled\": true }"
----

'''

=== Deleting a specific keystore configuration

The following example shows how to delete a specific keystore
configuration. Only an inactive configuration can be deleted.

----

# The API:
DELETE /api/security/key-stores/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-stores/33421d82-0a8d-
11ec-ae88-005056bb5955' -H 'accept: application/hal+json'
----

'''





[[IDd10ddb587e5d25e9b6aa8ebea779007f]]
= Retrieve keystores

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-stores`#

*Introduced In:* 9.10

Retrieves keystores.

== Related ONTAP commands
```

```
* `security key-manager show-key-store`
* `security key-manager keystore show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|configuration.uuid
|string
|query
|False
a|Filter by configuration.uuid

* Introduced in: 9.14


|configuration.name
|string
|query
|False
a|Filter by configuration.name

* Introduced in: 9.14


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|state
|string
|query
|False
a|Filter by state

* Introduced in: 9.16


|enabled
|boolean
|query
|False
a|Filter by enabled

* Introduced in: 9.14


|type
|string
|query
|False
a|Filter by type


|location
|string
|query
|False
a|Filter by location


|uuid
|string
|query
|False
a|Filter by uuid


|scope
|string
|query
|False
a|Filter by scope

* Introduced in: 9.14
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#security_keystore[security_keystore]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "configuration": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
      },
      "location": "string",
      "scope": "string",
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13434920
| SVM does not exist.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#security_keystore]
[.api-collapsible-fifth-title]
security_keystore

[cols=3*,options=header]
|===
|Name
|Type
|Description

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|location
|string
a|Indicates whether the keystore is onboard or external. * 'onboard' -
Onboard Key Database * 'external' - External Key Database, including KMIP
and Cloud Key Management Systems


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".
```

```
|state
|string
a|State of the keystore: * 'active' - The key manager is active and
serving new and existing keys. * 'mixed' - The key manager has a mixed
configuration. New keys can't be created. * 'svm_kek_rekey' - An SVM key
encryption key (KEK) rekey is in progress. New keys can't be created. *
'blocked' - The key manager is blocked and cannot serve new and existing
keys. * 'switching' - Switching the enabled key manager keystore
configuration. Some operations are blocked. * 'initializing' - The key
manager is being initialized. All operations are blocked. * 'disabling' -
The key manager is being disabled. All operations are blocked.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of keystore that is configured: * 'okm' - Onboard Key Manager *
'kmip' - External Key Manager * 'akv' - Azure Key Vault Key Management
Service * 'gcp' - Google Cloud Platform Key Management Service * 'aws' -
Amazon Web Service Key Management Service * 'ikp' - IBM Key Protect Key
Management Service * 'barbican' - Barbican Key Management Service


|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID0c6609827b17183b57986bc59ded4098]]
= Delete an inactive keystore configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/key-stores/{uuid}`#

*Introduced In:* 9.14

Deletes an inactive keystore configuration.

== Related ONTAP commands

* `security key-manager keystore delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 262155
| This operation requires an effective cluster version of 9.14.0 or later.

| 65538905
| The keystore configuration is currently enabled and cannot be deleted.

| 65538907
| The method is not yet supported for deleting the given UUID's type of
configuration.

| 65538908
| The specified keystore configuration UUID either does not exist or
corresponds to a keystore configuration that is not supported by this
operation.

| 65539521
| An effective cluster version of ONTAP 9.16.1 or later is required to
delete an inactive key manager on the admin SVM.

| 65539522
| This command does not support disabling key manager configurations with
the specified keystore type on the admin SVM.

| 65539593
| Cannot delete an inactive key manager configuration while the keystore
is in the switching state.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID18ba1971a1cdb9d285f8b04336be16f1]]
= Retrieve keystore configuration details with the specified UUID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/key-stores/{uuid}`#
```

```
*Introduced In:* 9.14

Retrieves details of the keystore configuration with the specified UUID.

== Related ONTAP commands

* `security key-manager keystore show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Keystore configuration UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|location
|string
a|Indicates whether the keystore is onboard or external. * 'onboard' -
Onboard Key Database * 'external' - External Key Database, including KMIP
and Cloud Key Management Systems


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|state
|string
a|State of the keystore: * 'active' - The key manager is active and
serving new and existing keys. * 'mixed' - The key manager has a mixed
configuration. New keys can't be created. * 'svm_kek_rekey' - An SVM key
encryption key (KEK) rekey is in progress. New keys can't be created. *
'blocked' - The key manager is blocked and cannot serve new and existing
keys. * 'switching' - Switching the enabled key manager keystore
configuration. Some operations are blocked. * 'initializing' - The key
manager is being initialized. All operations are blocked. * 'disabling' -
The key manager is being disabled. All operations are blocked.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of keystore that is configured: * 'okm' - Onboard Key Manager *
'kmip' - External Key Manager * 'akv' - Azure Key Vault Key Management
Service * 'gcp' - Google Cloud Platform Key Management Service * 'aws' -
Amazon Web Service Key Management Service * 'ikp' - IBM Key Protect Key
```

```
Management Service * 'barbican' - Barbican Key Management Service


|uuid
|string
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "configuration": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "location": "string",
  "scope": "string",
  "state": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 65539523
| Internal error. The keymanager_config table is missing.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID66e7b88560af13c2e74b61b9c5bd76b8]]
= Enable a keystore configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/key-stores/{uuid}`#

*Introduced In:* 9.14

Enables a keystore configuration

== Related ONTAP commands

* `security key-manager keystore enable`
* `security key-manager keystore disable`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|Keystore configuration UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|location
```

```
|string
a|Indicates whether the keystore is onboard or external. * 'onboard' -
Onboard Key Database * 'external' - External Key Database, including KMIP
and Cloud Key Management Systems


|state
|string
a|State of the keystore: * 'active' - The key manager is active and
serving new and existing keys. * 'mixed' - The key manager has a mixed
configuration. New keys can't be created. * 'svm_kek_rekey' - An SVM key
encryption key (KEK) rekey is in progress. New keys can't be created. *
'blocked' - The key manager is blocked and cannot serve new and existing
keys. * 'switching' - Switching the enabled key manager keystore
configuration. Some operations are blocked. * 'initializing' - The key
manager is being initialized. All operations are blocked. * 'disabling' -
The key manager is being disabled. All operations are blocked.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Type of keystore that is configured: * 'okm' - Onboard Key Manager *
'kmip' - External Key Manager * 'akv' - Azure Key Vault Key Management
Service * 'gcp' - Google Cloud Platform Key Management Service * 'aws' -
Amazon Web Service Key Management Service * 'ikp' - IBM Key Protect Key
Management Service * 'barbican' - Barbican Key Management Service


|uuid
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "configuration": {
    "name": "default",
```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
  },
  "location": "string",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262155
| This operation requires an effective cluster version of 9.14.0 or later.

| 65536203
| Internal error. Failed to generate SVM key. Cluster key database is not
initialized.

| 65536204
| Internal error. Failed to generate SVM key. Cluster key encryption key
is not found.

| 65536207
| Internal error. Failed to store SVM key.

| 65536212
| Internal error. Failed to delete SVM volume key.

| 65536217
| Internal error. Failed to delete key from cryptomod.

| 65536602
| Cannot perform enable/switch while nodes are out of quorum.

| 65536856
| No external key management server configured for SVM.

| 65536882
| Internal error. UUID is missing for the volume.

| 65536883
| Internal error. Volume encryption key is missing.

| 65536884
| Internal error. Volume encryption key is invalid for the volume.

| 65536889
```

```
| Internal error. The volume has an invalid encryption blob.

| 65536890
| Internal error. Failed to generate VDEK blob.

| 65536973
| Internal error. Volume DSID is missing for the volume.

| 65536987
| One or more key servers are unavailable.

| 65537527
| Internal error. Failed to rewrap SVM key encryption key.

| 65537533
| Internal error. Failed to unwrap SVM key encryption key.

| 65537605
| Failed to establish connectivity with the cloud key management service.

| 65538908
| The specified keystore configuration UUID either does not exist or
corresponds to a keystore configuration that is not supported by this
operation.

| 65538909
| A value for enabled is required.

| 65538910
| Disabling an enabled configuration through this method is currently not
supported.

| 65539200
| The key custodian was not provided with an external key value store.

| 65539201
| Failed to encrypt.

| 65539205
| This command does not support enabling key manager configurations with
the specified keystore type.

| 65539206
| The SVM associated with the supplied keystore UUID already has a
keystore configuration enabled. This command does not support migrating
from configurations of that keystore type".
```

```
| 65539207
| The specified keystore configuration does not exist.

| 65539212
| Cannot switch the enabled keystore configuration when it is not in the
'active' or 'switching' state.

| 65539218
| The specified keystore configuration is already enabled.

| 65539222
| Switching between configurations of different keystore types is not
supported for data SVMs.

| 65539416
| Internal error. Failed to parse key value store response.

| 65539442
| Internal error. Failed to generate AES blob.

| 65539444
| Internal error. Failed to update the VDEK blob.

| 65539513
| An effective cluster version of ONTAP 9.16.1 or later is required to
enable an inactive key manager on the admin SVM.

| 65539514
| This command does not support enabling key manager configurations with
the specified keystore type on the admin SVM.

| 65539515
| Cannot switch keystore types on the admin SVM. The keystore type for the
invalid configuration must be OKM and the enabled configuration must be
KMIP, or vice versa.

| 65539518
| Internal error. Cannot find the enabled configuration.

| 65539520
| Cannot enable the Onboard Key Manager on the admin SVM because an
inactive Onboard Key Manager already exists on the admin SVM.

| 65539534
| Cannot switch admin SVM Key Manager when system root volumes are
present.
```

```
| 65539535
| Failed to find unwrapped key on any nodes.

| 65539536
| Internal error. Failed to find the existing wrapped key blob for the
volume.

| 65539538
| SVM-KEK not created for the SVM.

| 65539539
| SVM-KEK not restored for SVM.

| 65539583
| Cannot switch to the Onboard Key Manager when the external key manager
has a policy associated with it

| 65539585
| Cannot enable an external key manager on the admin SVM because an
inactive external key manager already exists on the admin SVM.

| 65539590
| Cannot switch to the Onboard Key Manager if there are more than two NSE-
AKs in the cluster.

| 65539591
| Cannot switch to the Onboard Key Manager if there are fewer than two
NSE-AKs in the cluster.

| 65539622
| CKEK not restored for SVM.

| 65539704
| The key manager cannot be configured because the SVM has NAE volumes.

| 65539837
| The key manager configured does not have any volumes to migrate.

| 65539838
| Cannot migrate to a cloud key manager when the external key manager has
a policy associated with it.

| 65539839
| Cannot migrate to a cloud key manager when the key manager is not in a
mixed state.

| 65539840
```

```
| Failed to clean up the keys stored on the external key servers.

| 65539842
| Failed to restore the SVM-KEK.

| 65539845
| Cannot migrate SVM volumes to the Onboard Key Manager when the key
manager is not in the mixed or active state.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#configuration]
[.api-collapsible-fifth-title]
configuration

Security keystore object reference.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the configuration.


|uuid
|string
a|Keystore UUID.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#security_keystore]
[.api-collapsible-fifth-title]
security_keystore

[cols=3*,options=header]
|===
|Name
|Type
|Description

|configuration
|link:#configuration[configuration]
a|Security keystore object reference.


|enabled
|boolean
a|Indicates whether the configuration is enabled.


|location
|string
a|Indicates whether the keystore is onboard or external. * 'onboard' -
Onboard Key Database * 'external' - External Key Database, including KMIP
and Cloud Key Management Systems


|state
|string
a|State of the keystore: * 'active' - The key manager is active and
serving new and existing keys. * 'mixed' - The key manager has a mixed
configuration. New keys can't be created. * 'svm_kek_rekey' - An SVM key
encryption key (KEK) rekey is in progress. New keys can't be created. *
'blocked' - The key manager is blocked and cannot serve new and existing
keys. * 'switching' - Switching the enabled key manager keystore
configuration. Some operations are blocked. * 'initializing' - The key
manager is being initialized. All operations are blocked. * 'disabling' -
The key manager is being disabled. All operations are blocked.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
```

```
|string
a|Type of keystore that is configured: * 'okm' - Onboard Key Manager *
'kmip' - External Key Manager * 'akv' - Azure Key Vault Key Management
Service * 'gcp' - Google Cloud Platform Key Management Service * 'aws' -
Amazon Web Service Key Management Service * 'ikp' - IBM Key Protect Key
Management Service * 'barbican' - Barbican Key Management Service


|uuid
|string
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage security login messages

:leveloffset: +1


[[ID81f1075058e248a5728ed79d645ffe66]]
= Manage security login messages



== Overview

You can use this API to display and manage the login messages
configuration. The GET request retrieves all of the login messages in the
cluster. GET operations on /security/login/messages/\{uuid} retrieve the
login messages configuration by UUID. PATCH operations on
/security/login/messages/\{uuid} update the login messages configuration
by UUID.
 +

'''

== Examples

=== Retrieving all of the login messages in the cluster

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-ip>/api/security/login/messages?fields=*" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
    "scope": "cluster",
```

```
    "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
****\n",
    "message": "#### Welcome to Cluster X ####\n",
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
      }
    }
  },
  {
    "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
    "scope": "svm",
    "svm": {
      "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
      "name": "svm1"
    },
    "message": "#### Welcome to SVM1 ####\n",
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-
005056bbef18"
      }
    }
  },
  {
    "uuid": "8ddee11e-a58c-11e8-85e0-005056bbef18",
    "scope": "svm",
    "svm": {
      "uuid": "8ddee11e-a58c-11e8-85e0-005056bbef18",
      "name": "svm3"
    },
    "banner": "*** WARNING: This system is for the use of authorized users
only. ****\n",
    "_links": {
      "self": {
        "href": "/api/security/login/messages/8ddee11e-a58c-11e8-85e0-
005056bbef18"
      }
    }
  },
  {
    "uuid": "f7e41c99-9ffa-11e8-a5dd-005056bbef18",
    "scope": "svm",
```

```
    "svm": {
      "uuid": "f7e41c99-9ffa-11e8-a5dd-005056bbef18",
      "name": "svm2"
    },
    "_links": {
      "self": {
        "href": "/api/security/login/messages/f7e41c99-9ffa-11e8-a5dd-
005056bbef18"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/security/login/messages?fields=*"
  }
}
}
----

'''

=== Retrieving the login messages configuration at the cluster scope

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?scope=cluster&fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
    "scope": "cluster",
    "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
****\n",
    "message": "#### Welcome to Cluster X ####\n",
```

```
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?scope=cluster&fields=*"
  }
}
}
----

'''

=== Retrieving the login banner configured at the cluster scope

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?scope=cluster&fields=banner" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
    "scope": "cluster",
    "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
****\n",
    "_links": {
      "self": {
        "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
```

```
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?scope=cluster&fields=banner"
  }
}
}
----

'''

=== Retrieving the login messages configuration of a specific SVM

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?svm.name=svm1&fields=*" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
    "scope": "svm",
    "svm": {
      "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
      "name": "svm1"
    },
    "message": "#### Welcome to SVM1 ####\n",
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-
005056bbef18"
      }
```

```
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?svm.name=svm1&fields=*"
  }
}
}
----

'''

=== Retrieving the login messages configuration by UUID, including all
fields

'''

----

# The API:
/api/security/login/messages/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/security/login/messages/7b1b3715-9ffa-
11e8-a5dd-005056bbef18?fields=*" -H "accept: application/hal+json"

# The response:
{
"uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
"scope": "svm",
"svm": {
  "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
  "name": "svm1"
},
"message": "#### Welcome to SVM1 ####\n",
"show_cluster_message": true,
"_links": {
  "self": {
    "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-
005056bbef18"
  }
}
}
----
```

```
'''

=== Configuring the login banner in a cluster

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"banner\": \"You are entering secure area.\" }"

# The response:
{
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?scope=cluster"
  }
}
}
----

'''

=== Configuring the message of the day (MOTD) in a cluster

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"message\": \"Welcome to Cluster X\",  \"show_cluster_message\": true }"

# The response:
{
```

```
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?scope=cluster"
  }
}
}
----

'''

=== Clearing the login banner and message of the day (MOTD) in a cluster

'''

----

# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"banner\": \"\", \"message\": \"\" }"

# The response:
{
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?scope=cluster"
  }
}
}
----

'''

=== Configuring the login messages for a specific SVM

'''

----

# The API:
/api/security/login/messages
```

```
# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?svm.name=svm1" -H  "accept:
application/hal+json" -H  "Content-Type: application/json" -d "{
\"banner\" : \"AUTHORIZED ACCESS ONLY\", \"message\": \"WELCOME!\" }"

# The response:
{
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/messages?svm.name=svm1"
  }
}
}
----

'''

=== Configuring the login messages by UUID

'''

----

# The API:
/api/security/login/messages/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/login/messages/7b1b3715-
9ffa-11e8-a5dd-005056bbef18" -H  "accept: application/hal+json" -H
"Content-Type: application/json" -d "{ \"banner\" : \"AUTHORIZED ACCESS
ONLY\", \"message\": \"WELCOME!\" }"
----

'''

=== Clearing the login messages configuration by UUID

'''

----

# The API:
/api/security/login/messages/{uuid}
```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/login/messages/7b1b3715-
9ffa-11e8-a5dd-005056bbef18" -H "accept: application/hal+json" -H
"Content-Type: application/json" -d "{ \"banner\": \"\", \"message\": \"\"
}"
----

'''





[[ID9f60160b9d791e8e4453b5f3ab73de0c]]
= Retrieve the login banner and messages of the day

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/login/messages`#

*Introduced In:* 9.6

Retrieves the login banner and messages of the day (MOTD) configured in
the cluster
and in specific SVMs.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|banner
|string
|query
|False
a|Filter by banner

* maxLength: 2048
* minLength: 0
```

```
|uuid
|string
|query
|False
a|Filter by uuid


|scope
|string
|query
|False
a|Filter by scope


|message
|string
|query
|False
a|Filter by message

* maxLength: 2048
* minLength: 0


|show_cluster_message
|boolean
|query
|False
a|Filter by show_cluster_message


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|fields
|array[string]
```

```
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#login_messages[login_messages]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "banner": "string",
```

```
      "message": "string",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
```

```
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#login_messages]
[.api-collapsible-fifth-title]
login_messages

The login banner and message of the day (MOTD) configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|banner
|string
a|The login banner text. This message is displayed during SSH and console
device
login just before the password prompt displays. When configured, a
cluster-level
login banner is used for every incoming connection. Each data SVM can
override
the cluster-level banner to instead display when you log into the SVM. To
restore
the default setting for a data SVM, set the banner to an empty string.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|message
|string
a|The message of the day (MOTD). This message appears just before the
clustershell
prompt after a successful login. When configured, the cluster message
displays first. If you log in as a data SVM administrator, the
SVM message is then printed. The cluster-level MOTD can be disabled
for a given data SVM using the "show_cluster_message" property.
```

```
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|show_cluster_message
|boolean
a|Specifies whether to show a cluster-level message before the SVM message
when logging in as an SVM administrator.
This setting can only be modified by cluster administrators.
Optional in the PATCH body.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|The unique identifier (ID) of the login messages configuration.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0f30ee29c09e7646504593e16fb2bb28]]
= Retrieve a login message configuration by UUID
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/login/messages/{uuid}`#

*Introduced In:* 9.6

Retrieves the login messages configuration by UUID.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Login messages configuration UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|banner
|string
a|The login banner text. This message is displayed during SSH and console
device
login just before the password prompt displays. When configured, a
cluster-level
login banner is used for every incoming connection. Each data SVM can
override
the cluster-level banner to instead display when you log into the SVM. To
restore
the default setting for a data SVM, set the banner to an empty string.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|message
|string
a|The message of the day (MOTD). This message appears just before the
clustershell
prompt after a successful login. When configured, the cluster message
displays first. If you log in as a data SVM administrator, the
SVM message is then printed. The cluster-level MOTD can be disabled
for a given data SVM using the "show_cluster_message" property.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|show_cluster_message
|boolean
a|Specifies whether to show a cluster-level message before the SVM message
when logging in as an SVM administrator.
This setting can only be modified by cluster administrators.
Optional in the PATCH body.


|svm
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|The unique identifier (ID) of the login messages configuration.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "banner": "string",
  "message": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
```

```
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDeb6ef894c50435428fd1e835ce94da64]]
= Update a login message configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/login/messages/{uuid}`#

*Introduced In:* 9.6

Updates the login messages configuration.
There are no required fields. An empty body makes no modifications.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
```

```
a|Login messages configuration UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|banner
|string
a|The login banner text. This message is displayed during SSH and console
device
login just before the password prompt displays. When configured, a
cluster-level
login banner is used for every incoming connection. Each data SVM can
override
the cluster-level banner to instead display when you log into the SVM. To
restore
the default setting for a data SVM, set the banner to an empty string.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|message
|string
a|The message of the day (MOTD). This message appears just before the
clustershell
prompt after a successful login. When configured, the cluster message
displays first. If you log in as a data SVM administrator, the
SVM message is then printed. The cluster-level MOTD can be disabled
for a given data SVM using the "show_cluster_message" property.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".
```

```
|show_cluster_message
|boolean
a|Specifies whether to show a cluster-level message before the SVM message
when logging in as an SVM administrator.
This setting can only be modified by cluster administrators.
Optional in the PATCH body.


|uuid
|string
a|The unique identifier (ID) of the login messages configuration.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "banner": "string",
  "message": "string",
  "scope": "string",
  "uuid": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error codes | Description

| 10225636
| Only a cluster administrator can modify the `show_cluster_message`
property.
|===
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#login_messages]
[.api-collapsible-fifth-title]
login_messages

The login banner and message of the day (MOTD) configuration.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|banner
|string
a|The login banner text. This message is displayed during SSH and console
device
login just before the password prompt displays. When configured, a
cluster-level
login banner is used for every incoming connection. Each data SVM can
override
the cluster-level banner to instead display when you log into the SVM. To
restore
the default setting for a data SVM, set the banner to an empty string.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.


|message
|string
a|The message of the day (MOTD). This message appears just before the
clustershell
prompt after a successful login. When configured, the cluster message
displays first. If you log in as a data SVM administrator, the
SVM message is then printed. The cluster-level MOTD can be disabled
for a given data SVM using the "show_cluster_message" property.
New lines are supplied as either LF or CRLF but are always returned as LF.
Optional in the PATCH body.
```

```
|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".


|show_cluster_message
|boolean
a|Specifies whether to show a cluster-level message before the SVM message
when logging in as an SVM administrator.
This setting can only be modified by cluster administrators.
Optional in the PATCH body.


|uuid
|string
a|The unique identifier (ID) of the login messages configuration.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and create TOTP profiles

:leveloffset: +1


[[IDe54146d024392ec90b4e42aa7ecc2f71]]
```

```
= View and create TOTP profiles



== Overview

This API configures the TOTP(Time based OTP) profile for user accounts for
secure shell (SSH) access.
Prerequisites:
  You must configure TOTP as a secondary authentication method in
/security/accounts when creating users.
  You must have a TOTP application that you can use to get TOTP keys.
  You are only allowed to create a TOTP profile for yourself, not for
others.
  Admin cannot create TOTP profiles for any users.

== Examples

=== Creating a TOTP profile for cluster-scoped user accounts

Specify the user account name and comment in the body of the POST request.
The owner.uuid or owner.name are not required for a cluster-scoped user
account.

----

# The API:
POST "/api/security/login/totps"

# The call:
curl -k https://<mgmt-ip>/api/security/login/totps --request POST --data
'{ "account": { "name": "pubuser2" },"comment": "Cserver-Creation"}'

# The response:
{
"num_records": 1,
"records": [
{
  "owner": {
    "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
```

```
  "account": {
    "name": "pubuser2"
  },
  "secret_key": "<ACCOUNT-TOTP-SECRET-KEY>"
  "install_url":
"https://www.google.com/chart?chs=200x200&chld=M&#124;0&cht=qr&chl=otpauth
://totp/root@node1%3Fsecret%3DDRY5CAJGTQCL5TV4D3UAMYXJFM%26issuer%3Dnode1"
,
  "verification_code": "946090",
  "emergency_codes": [
    "54200192",
    "10418385",
    "52726505",
    "41704451",
    "20744310"
  ],
  "_links": {
    "self": {
      "href": "/api/security/login/totps/b009a9e7-4081-b576-7575-
ada21efcaf16/pubuser2"
    }
  }
}
 ]
}
----

=== Creating a TOTP profile for SVM-scoped user accounts

For an SVM-scoped account, specify either the SVM name as the owner.name
or the SVM UUID as the owner.uuid along with other parameters for the user
account. These parameters indicate the SVM that contains the user account
for the TOTP profile being created and can be obtained from the response
body of the GET request performed on the API "/api/svm/svms".

----

# The API:
POST "/api/security/login/totps"

# The call:
curl -k https://<mgmt-ip>/api/security/login/totps --request POST --data
'{ "account": { "name": "pubuser4" },"comment": "Vserver-
Creation","owner.name":"vs0" }'

# The response:
```

```
{
"num_records": 1,
"records": [
{
  "owner": {
    "uuid": "b019a9e7-4081-b576-7575-ada21efcaf16",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b019a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
  "account": {
    "name": "pubuser4"
  },
  "secret_key": "<ACCOUNT-TOTP-SECRET-KEY>",
  "install_url":
"https://www.google.com/chart?chs=200x200&chld=M&#124;0&cht=qr&chl=otpauth
://totp/root@node1%3Fsecret%3DDRY5CAJGTQCL5TV4D3UAMYXJFM%26issuer%3Dnode1"
,
  "verification_code": "946090",
  "emergency_codes": [
    "54200192",
    "10418385",
    "52726505",
    "41704451",
    "20744310"
  ],
  "_links": {
    "self": {
      "href": "/api/security/login/totps/b019a9e7-4081-b576-7575-
ada21efcaf16/pubuser4"
    }
  }
}
 ]
}
----

=== Retrieving the configured TOTP profile for user accounts

Retrieves the TOTP profiles associated with the user accounts or a
filtered list (for a specific user account name, a specific SVM and so
on).
```

```
----

# The API:
GET "/api/security/login/totps"

# The call to retrieve TOTP profiles associated with TOTP configured user
accounts in the cluster:
curl -k https://<mgmt-ip>/api/security/login/totps

# the response:
{
"records": [
{
  "owner": {
    "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
  "account": {
    "name": "pubuser2",
    "_links": {
      "self": {
        "href": "/api/security/accounts/b009a9e7-4081-b576-7575-
ada21efcaf16/pubuser2"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/security/login/totps/b009a9e7-4081-b576-7575-
ada21efcaf16/pubuser2"
    }
  }
}
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/login/totps"
  }
}
}
```

```
----




[[ID3e3c4a4aa5c02967ef0655e4bb6f0018]]
= Retrieve TOTP profiles configured for user accounts

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/login/totps`#

*Introduced In:* 9.13

Retrieves the TOTP profiles configured for user accounts.

== Related ONTAP commands

* `security login totp show`

== Learn more

* xref:{relative_path}security_login_totps_endpoint_overview.html[DOC
/security/login/totps]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|account.name
|string
|query
|False
a|Filter by account.name


|owner.name
```

```
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|comment
|string
|query
|False
a|Filter by comment


|enabled
|boolean
|query
|False
a|Filter by enabled


|sha_fingerprint
|string
|query
|False
a|Filter by sha_fingerprint


|scope
|string
|query
|False
a|Filter by scope


|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#totp[totp]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "account": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
          }
        },
        "name": "joe.smith"
      },
      "comment": "string",
      "enabled": "",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "string",
      "sha_fingerprint": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account
```

```
|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the TOTP profile.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#totp]
[.api-collapsible-fifth-title]
totp

TOTP profile for the user account used to access SSH.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|account
|link:#account_reference[account_reference]
a|

|comment
|string
a|Optional comment for the TOTP profile.


|enabled
|boolean
a|Status of the TOTP profile.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[IDcb16d7a218c65bce62e91e056543b361]]
= Create the TOTP profile for a user account

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/login/totps`#

*Introduced In:* 9.13

Creates a TOTP profile for a user account.

== Required properties

* `owner.uuid` - Account owner UUID.
* `account.name` - Account user name.

== Related ONTAP commands

* `security login totp create`

== Learn more

* xref:{relative_path}security_login_totps_endpoint_overview.html[DOC
/security/login/totps]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.
```

```
* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|

|comment
|string
a|Optional comment for the TOTP profile.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "account": {
    "name": "joe.smith"
  },
  "comment": "string",
  "enabled": "",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "sha_fingerprint": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#totp_post[totp_post]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "account": {
```

```
        "name": "joe.smith"
      },
      "comment": "string",
      "emergency_codes": "17503785",
      "enabled": 1,
      "install_url": "string",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "string",
      "secret_key": "string",
      "sha_fingerprint": "string",
      "verification_code": "string"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```


Status: Default


```
ONTAP Error Response Codes

|===
| Error Code | Description
```

```
| 144834561
| TOTP is not configured for the user.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the TOTP profile.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#totp]
[.api-collapsible-fifth-title]
totp

TOTP profile for the user account used to access SSH.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|account
|link:#account_reference[account_reference]
a|

|comment
|string
a|Optional comment for the TOTP profile.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


[#totp_post]
[.api-collapsible-fifth-title]
totp_post

Response object of the TOTP profile creation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|account
|link:#account_reference[account_reference]
a|

|comment
|string
```

```
a|Optional comment for the TOTP profile.


|emergency_codes
|array[string]
a|TOTP profile emergency codes for a user. These codes are for emergency
use when a user cannot access 2FA codes through other means.


|install_url
|string
a|TOTP profile installation URL for a user.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|secret_key
|string
a|TOTP profile secret key for a user.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|verification_code
|string
a|TOTP profile verification code for a user.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage TOTP profiles

:leveloffset: +1


[[IDdff562ec019de8ca81739180a24bb03f]]
= Manage TOTP profiles



== Overview

This API configures the TOTP profile for user accounts.
Specify the owner UUID and the account user name. The owner UUID
corresponds to the UUID of the SVM containing the user account associated
with the TOTP profile and can be obtained from the response body of the
GET request performed on the API "`/api/svm/svms".

== Examples

=== Retrieving the specific configured TOTP profile for user accounts

----

# The API:
GET "/api/security/login/totps/{owner.uuid}/{account.name}"

# The call:
curl -k https://<mgmt-ip>/api/security/login/totps/513a78c7-8c13-11e9-
8f78-005056bbf6ac/pubuser4

# the response:
{
"owner": {
```

```
    "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "Default",
  "_links": {
    "self": {
      "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
    }
  }
},
"account": {
  "name": "pubuser2",
  "_links": {
    "self": {
      "href": "/api/security/accounts/b009a9e7-4081-b576-7575-
ada21efcaf16/pubuser2"
    }
  }
},
"sha_fingerprint":
"21364f5417600e3d9d6a7ac6c05dd244aed9f15dce6786a2c89399a41ff0fdb0",
"scope": "cluster",
"_links": {
  "self": {
    "href": "/api/security/login/totps/b009a9e7-4081-b576-7575-
ada21efcaf16/pubuser2"
  }
}
}
----

=== Modifying the TOTP profile for a user account

----

# The API:
PATCH "/api/security/login/totps/{owner.uuid}/{account.name}"

# The call:
curl -k "https://<mgmt-ip>/api/security/login/totps/6865196a-8b59-11ed-
874c-0050568e36ed/ysadmin" --request PATCH --data "{  \"comment\":
\"Testing\",  \"enabled\": false}"

# the response:
{}
----

=== Deleting the TOTP profile for user accounts
```

```
----

# The API:
DELETE "/api/security/login/totps/{owner.uuid}/{account.name}"

# The call:
curl -k https://<mgmt-ip>/api/security/login/totps/d49de271-8c11-11e9-
8f78-005056bbf6ac/pubuser1 --request DELETE

# the response:
{}
----





[[IDbfb57b827e4fe78da61071a6d2a7040c]]
= Delete the TOTP profile for a user account

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/login/totps/{owner.uuid}/{account.name}`#

*Introduced In:* 9.13

Deletes the TOTP profile for a user account.

== Related ONTAP commands

* `security login totp delete`

== Learn more

*
xref:{relative_path}security_login_totps_owner.uuid_account.name_endpoint_
overview.html[DOC /security/login/totps/{owner.uuid}/{account.name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.


|account.name
|string
|path
|True
a|Account user name.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID13ae3f4f074a344ca33f78fab0f254e4]]
= Retrieve the TOTP profile for a user account

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/login/totps/{owner.uuid}/{account.name}`#

*Introduced In:* 9.13

Retrieves the TOTP profile configured for a user account.

== Related ONTAP commands
```

```
* `security login totp show`

== Learn more

*
xref:{relative_path}security_login_totps_owner.uuid_account.name_endpoint_
overview.html[DOC /security/login/totps/{owner.uuid}/{account.name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID.


|account.name
|string
|path
|True
a|Account user name.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|account
|link:#account_reference[account_reference]
a|

|comment
|string
a|Optional comment for the TOTP profile.


|enabled
|boolean
a|Status of the TOTP profile.


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


.Example response
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "account": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "joe.smith"
  },
  "comment": "string",
  "enabled": "",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "string",
  "sha_fingerprint": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the TOTP profile.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDe2d58aa207f771afec9c6b379984aa5d]]
= Update a TOTP profile for a user account

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/login/totps/{owner.uuid}/{account.name}`#

*Introduced In:* 9.13

Updates a TOTP user account.
```

```
== Related ONTAP commands

* `security login totp modify`

== Learn more

*
xref:{relative_path}security_login_totps_owner.uuid_account.name_endpoint_
overview.html[DOC /security/login/totps/{owner.uuid}/{account.name}]
* xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Account owner UUID


|account.name
|string
|path
|True
a|User account name

|===

== Request Body



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|comment
|string
a|Optional comment for the TOTP profile.


|enabled
|boolean
a|Status of the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "enabled": "",
  "scope": "string",
  "sha_fingerprint": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 144834564
| Only users with the admin role are allowed to modify the TOTP status.

| 144834565
| Invalid option for the field -enabled

| 144834566
| The user does not have a TOTP configuration available for modification.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#account_reference]
[.api-collapsible-fifth-title]
account_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|User account


|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the TOTP profile.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#totp]
[.api-collapsible-fifth-title]
totp

TOTP profile for the user account used to access SSH.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Optional comment for the TOTP profile.


|enabled
|boolean
a|Status of the TOTP profile.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|sha_fingerprint
|string
a|SHA fingerprint for the TOTP secret key.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage the multi-admin-verify global setting

:leveloffset: +1


[[IDbf8781271c7e9e04e58f6c56ec38b54d]]
= Manage the multi-admin-verify global setting



== Overview

These APIs provide information on the multi-admin verification global
setting.
The GET API retrieves the object store that contains the global setting
values of the multi-admin-verify feature.
The PATCH request is used to modify the multi-admin-verify global setting.
All fields are optional for the PATCH request.
Note it is recommended that multi-admin-verify is enabled equally on
peered ONTAP clusters.
 +

'''

== Examples

=== Retrieving the multi-admin-verify global setting

Retrieves the current multi-admin-verify global setting. If the global
setting is not set, default values are returned.
 +

'''

----

# The API:
/api/security/multi-admin-verify
```

```
# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-verify"

# The response:
{
"approval_groups": [
],
"required_approvers": 1,
"enabled": false,
"execution_expiry": "PT1H",
"approval_expiry": "PT1H",
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify"
  }
}
}
----

'''

=== Updating the multi-admin-verify global setting

The following example updates the multi-admin-verify global settings.
Note that the approval_groups needs to be available in /security/multi-
admin-verify/approval-groups before it is set in the global setting.
Note that the total number of approvers in an approval group must be a
positive integer more than the number of required approvers. For example,
if there are a total of 10 approvers in a group, the required approvers
can be set between 1 and 9.
 +

'''

----

# The API:
/api/security/multi-admin-verify

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/multi-admin-verify" -d
'{"required_approvers": "1", "enabled": "true", "execution_expiry": "2h",
"approval_expiry": "3h"}'
----

'''
```

```
[[ID266a9650224aa74051324782718c9435]]
= Retrieve a multi-admin-verify configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify`#

*Introduced In:* 9.11

Retrieves the multi-admin-verify configuration.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
```

```
a|Default time for requests to be approved, in ISO-8601 duration format.


|approval_groups
|array[string]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|enabled
|boolean
a|

|execution_expiry
|string
a|Default time for requests to be executed once approved, in ISO-8601
duration format.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approval_expiry": "string",
  "approval_groups": [
    "string"
  ],
  "execution_expiry": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



[[IDd15bcd2ab1cc1cecf67f9ee36d7eef10]]
= Modify a multi-admin-verify configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/multi-admin-verify`#

*Introduced In:* 9.11

Modifies the multi-admin-verify configuration.


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Default time for requests to be approved, in ISO-8601 duration format.


|approval_groups
|array[string]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|enabled
|boolean
a|

|execution_expiry
|string
a|Default time for requests to be executed once approved, in ISO-8601
duration format.
```

```
|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approval_expiry": "string",
  "approval_groups": [
    "string"
  ],
  "execution_expiry": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262309
| The feature must be enabled first.

| 262311
| Value must be greater than zero.

| 262312
| Number of required approvers must be less than the total number of
unique approvers in the approval-groups.
```

```
| 262313
| Number of unique approvers in the approval-groups must be greater than
the number of required approvers.

| 262314
| One or more approval groups added do not exist.

| 262315
| Approval-groups must be specified when enabling this feature.

| 262316
| Value must be in the range two minutes to two weeks.

| 262318
| multi-admin-verify requires an effective cluster version of ONTAP 9.11.1
or later.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#multi_admin_verify_config]
[.api-collapsible-fifth-title]
multi_admin_verify_config

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Default time for requests to be approved, in ISO-8601 duration format.


|approval_groups
```

```
|array[string]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|enabled
|boolean
a|

|execution_expiry
|string
a|Default time for requests to be executed once approved, in ISO-8601
duration format.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View and create multi-admin-verify approval groups

:leveloffset: +1


[[ID8ffbc19cfb8fa941ca5f9e3ec92e145]]
```

```
= View and create multi-admin-verify approval groups



== Overview

A group of users can be defined in a cluster server context.
Approval groups can be associated with a rule or global setting from which
the associated request can retrieve approvals.
 +

'''

== Examples

=== Creating a multi-admin-verify approval group

Creates an approval group for a specified SVM for a specified list of
ONTAP users.
 +

'''

----

# The API:
/api/security/multi-admin-verify/approval-groups

# The call:
curl -X POST "https://<mgmt-ip>/api/security/multi-admin-verify/approval-
groups?return_records=true" -H "accept: application/hal+json" -d '{"name":
"group1", "approvers": ["admin"], "email":
["group1.approvers@email.com"]}'

# The response:
{
"num_records": 1,
"records": [
  {
    "owner": {
      "uuid": "c109634f-7011-11ec-a23d-005056a78fd5",
      "_links": {
        "self": {
          "href": "/api/svm/svms/c109634f-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
```

```
    "name": "group1",
    "approvers": [
      "admin"
    ],
    "email": [
      "group1.approvers@email.com"
    ],
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/approval-
groups/c109634f-7011-11ec-a23d-005056a78fd5/group1"
      }
    }
  }
]
}
----

'''

=== Retrieving multi-admin-verify approval groups

Displays information about approval groups and the users that are
registered with each group.
 +

'''

----

# The API:
/api/security/multi-admin-verify/approval-groups

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-
verify/approval-groups"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
```

```
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "name": "group1",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/approval-
groups/52b75787-7011-11ec-a23d-005056a78fd5/group1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/approval-groups"
  }
}
}
----

'''
```


```
[[ID169a181265ed21109ed74e4823052cef]]
= Retrieve multi-admin-verify approval groups

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/approval-groups`#

*Introduced In:* 9.11

Retrieves multi-admin-verify approval groups.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|approvers
|string
|query
|False
a|Filter by approvers


|email
|string
|query
|False
a|Filter by email


|name
|string
|query
|False
a|Filter by name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|owner.name
|string
|query
|False
a|Filter by owner.name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
```

```
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_approval_group[multi_admin_verify_approval
_group]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "approvers": [
        "string"
      ],
      "email": [
        "string"
      ],
      "name": "string",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the approval group. The only valid owner is currently the
cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#multi_admin_verify_approval_group]
[.api-collapsible-fifth-title]
multi_admin_verify_approval_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approvers
|array[string]
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|name
|string
a|Name of the approval group.


|owner
|link:#owner[owner]
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID94667480b6895c5334cc15d193a45512]]
= Create a multi-admin-verify approval group

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/multi-admin-verify/approval-groups`#

*Introduced In:* 9.11

Creates a multi-admin-verify approval group.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approvers
|array[string]
```

```
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|name
|string
a|Name of the approval group.


|owner
|link:#owner[owner]
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approvers": [
    "string"
  ],
  "email": [
    "string"
  ],
  "name": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_approval_group[multi_admin_verify_approval
_group]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "approvers": [
        "string"
      ],
      "email": [
        "string"
      ],
      "name": "string",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
```

```
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262309
| The feature must be enabled first.

| 262333
| Invalid user in approval group.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the approval group. The only valid owner is currently the
cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#multi_admin_verify_approval_group]
```

```
[.api-collapsible-fifth-title]
multi_admin_verify_approval_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approvers
|array[string]
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|name
|string
a|Name of the approval group.


|owner
|link:#owner[owner]
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage multi-admin-verify approval groups

:leveloffset: +1


[[ID0c6406effdcc5586dfa5dbd5d321fbcc]]
= Manage multi-admin-verify approval groups



== Overview

These APIs provide information about a specific multi-admin verification
approval-group.
A group of users can be defined in a cluster server context.
Approval groups can be associated with a rule or global setting from which
the associated request can retrieve approvals.
 +

'''

== Examples

=== Retrieving a multi-admin-verify approval group

Displays information about a specific approval group and the users that
are registered within that group.
 +

'''

----

# The API:
/api/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-
verify/approval-groups/52b75787-7011-11ec-a23d-005056a78fd5/group1"

# The response:
{
"owner": {
```

```
  "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
  "name": "cluster1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
    }
  }
},
"name": "group1",
"approvers": [
  "admin"
],
"email": [
  "group1.approvers@email.com"
],
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/approval-groups/52b75787-
7011-11ec-a23d-005056a78fd5/group1"
  }
}
}
----

'''

=== Updating a multi-admin-verify approval group

Modifies attributes of an approval group.
 +

'''

----

# The API:
/api/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/multi-admin-
verify/approval-groups/52b75787-7011-11ec-a23d-005056a78fd5/group1" -d
'{"approvers": ["admin1"], "email": ["group1.approvers.new@email.com"]}'
----

'''
```

```
=== Deleting a multi-admin-verify approval group

Deletes the specified approval group.
 +

'''

----

# The API:
/api/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}

# The call:
curl -X DELETE "https://<cluster-ip>/api/security/multi-admin-
verify/approval-groups/52b75787-7011-11ec-a23d-005056a78fd5/group1"
----

'''





[[IDeff469cd70dfb0c30c217136521c858d]]
= Delete a multi-admin-verify approval group

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}`#

*Introduced In:* 9.11

Deletes a multi-admin-verify approval group.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
```

```
|path
|True
a|

|name
|string
|path
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262319
| Cannot delete an approval group that is in the list of global approval
groups.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDdca8c0d14129c233ba2a581398a402c8]]
= Retrieve a multi-admin-verify approval group
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}`#

*Introduced In:* 9.11

Retrieves a multi-admin-verify approval group.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|name
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|approvers
|array[string]
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|name
|string
a|Name of the approval group.


|owner
|link:#owner[owner]
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approvers": [
    "string"
  ],
  "email": [
    "string"
  ],
  "name": "string",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
```

```
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the approval group. The only valid owner is currently the
cluster.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0a2e7dffb8a0ca0255be4130bbbc8112]]
= Update a multi-admin-verify approval group

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/multi-admin-verify/approval-groups/{owner.uuid}/{name}`#

*Introduced In:* 9.11

Updates a multi-admin-verify approval group.
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|name
|string
|path
|True
a|
|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approvers
|array[string]
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|owner
|link:#owner[owner]
```

```
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approvers": [
    "string"
  ],
  "email": [
    "string"
  ],
  "owner": {
    "name": "svm1"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262331
| At least one approver is required.

| 262332
| An add or remove list is required.

| 262333
| Invalid user in approval group.
```

```
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the approval group. The only valid owner is currently the
cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|===


[#multi_admin_verify_approval_group]
[.api-collapsible-fifth-title]
multi_admin_verify_approval_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approvers
|array[string]
a|List of users that can approve a request.


|email
|array[string]
a|Email addresses that are notified when a request is created, approved,
vetoed, or executed.


|owner
|link:#owner[owner]
a|The owner of the approval group. The only valid owner is currently the
cluster.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage multi-admin-verify approval requests

:leveloffset: +1


[[IDb2e597da7fd88f32ffb9a94975de967d]]
= Manage multi-admin-verify approval requests



== Overview

These APIs provide information about multi-admin verification requests.
If you need to execute a command that is protected by a multi-admin rule,
you must first submit a request to be allowed to execute the command.
The request must then be approved by the designated approvers according to
the rule associated with the command.
 +

'''

== Examples

=== Creating a multi-admin-verify request

Creates a request for the specified ONTAP operation.
 +

'''

----

# The API:
/api/security/multi-admin-verify/requests

# The call:
curl -X POST "https://<mgmt-ip>/api/security/multi-admin-
```

```
verify/requests?return_records=true" -H "accept: application/hal+json" -d
'{"operation": "volume delete", "query": "-vserver vs0 -volume v1",
"permitted_users": ["user1","user2"]}'

# The response:
{
"num_records": 1,
"records": [
  {
    "index": 10,
    "operation": "volume delete",
    "query": "-vserver vs0 -volume v1",
    "state": "pending",
    "required_approvers": 2,
    "pending_approvers": 2,
    "execute_on_approval": false,
    "permitted_users": [
      "user1",
      "user2"
    ],
    "user_requested": "admin",
    "owner": {
      "uuid": "c1483186-6e73-11ec-bc92-005056a7ad04",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/c1483186-6e73-11ec-bc92-005056a7ad04"
        }
      }
    },
    "create_time": "2022-01-06T16:59:49-05:00",
    "approve_expiry_time": "2022-01-06T19:59:49-05:00",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/requests/10"
      }
    }
  }
]
}
----

'''

=== Retrieving multi-admin-verify requests
```

```
Retrieves information about multi-admin verification requests.
 +

'''

----

# The API:
/api/security/multi-admin-verify/requests

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-
verify/requests"

# The response:
{
"records": [
  {
    "index": 1,
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/requests/1"
      }
    }
  },
  {
    "index": 2,
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/requests/2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/requests"
  }
}
}
----

'''
```

```
[[ID38d45f522ab781340f6cb424171a1195]]
= Retrieve multi-admin-verify requests

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/requests`#

*Introduced In:* 9.11

Retrieves multi-admin-verify requests.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|approved_users
|string
|query
|False
a|Filter by approved_users


|index
|integer
|query
|False
a|Filter by index


|create_time
|string
|query
|False
a|Filter by create_time


|execution_expiry_time
|string
```

```
|query
|False
a|Filter by execution_expiry_time


|potential_approvers
|string
|query
|False
a|Filter by potential_approvers


|approve_time
|string
|query
|False
a|Filter by approve_time


|query
|string
|query
|False
a|Filter by query


|approve_expiry_time
|string
|query
|False
a|Filter by approve_expiry_time


|user_vetoed
|string
|query
|False
a|Filter by user_vetoed


|user_requested
|string
|query
|False
a|Filter by user_requested
```

```
|permitted_users
|string
|query
|False
a|Filter by permitted_users


|pending_approvers
|integer
|query
|False
a|Filter by pending_approvers


|comment
|string
|query
|False
a|Filter by comment


|owner.name
|string
|query
|False
a|Filter by owner.name


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|execute_on_approval
|boolean
|query
|False
a|Filter by execute_on_approval

* Introduced in: 9.13


|required_approvers
|integer
|query
```

```
|False
a|Filter by required_approvers


|operation
|string
|query
|False
a|Filter by operation


|state
|string
|query
|False
a|Filter by state


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
```

```
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_request[multi_admin_verify_request]]
a|

|===


.Example response
[%collapsible%closed]
```

```
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "approve_expiry_time": "string",
      "approve_time": "string",
      "approved_users": [
        "string"
      ],
      "comment": "string",
      "create_time": "string",
      "execution_expiry_time": "string",
      "index": 0,
      "operation": "string",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "pending_approvers": 0,
      "permitted_users": [
        "string"
      ],
      "potential_approvers": [
        "string"
      ],
      "query": "string",
      "required_approvers": 0,
      "state": "string",
      "user_requested": "string",
      "user_vetoed": "string"
    }
```

```
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the request. This can identify the cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#multi_admin_verify_request]
[.api-collapsible-fifth-title]
multi_admin_verify_request

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.


|comment
|string
a|Optional user-provided comment that is sent to the approval-group email
indicating why the request was made.


|create_time
|string
a|

|execute_on_approval
|boolean
a|Specifies that the operation is executed automatically on final
approval.


|execution_expiry_time
|string
a|

|index
|integer
a|Unique index that represents a request.


|operation
|string
a|The command to execute.


|owner
```

```
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|permitted_users
|array[string]
a|List of users that can execute the operation once approved. If not set,
any authorized user can perform the operation.


|potential_approvers
|array[string]
a|The users that are able to approve the request.


|query
|string
a|Identifies the specific entry upon which the user wants to operate.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
required.


|user_requested
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true
```

```
|user_vetoed
|string
a|The user that vetoed the request.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb4ab354d7ce10f966eb50411b4d6057f]]
= Create a multi-admin-verify request

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/multi-admin-verify/requests`#

*Introduced In:* 9.11

Creates a multi-admin-verify request.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
```

```
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.


|comment
|string
a|Optional user-provided comment that is sent to the approval-group email
indicating why the request was made.


|create_time
|string
a|

|execution_expiry_time
|string
a|

|index
|integer
a|Unique index that represents a request.
```

```
|operation
|string
a|The command to execute.


|owner
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|permitted_users
|array[string]
a|List of users that can execute the operation once approved. If not set,
any authorized user can perform the operation.


|potential_approvers
|array[string]
a|The users that are able to approve the request.


|query
|string
a|Identifies the specific entry upon which the user wants to operate.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
required.


|user_requested
```

```
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|user_vetoed
|string
a|The user that vetoed the request.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approve_expiry_time": "string",
  "approve_time": "string",
  "approved_users": [
    "string"
  ],
  "comment": "string",
  "create_time": "string",
  "execute_on_approval": true,
  "execution_expiry_time": "string",
  "index": 0,
  "operation": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "pending_approvers": 0,
  "permitted_users": [
    "string"
  ],
  "potential_approvers": [
    "string"
  ],
  "query": "string",
  "required_approvers": 0,
  "state": "string",
```

```
  "user_requested": "string",
  "user_vetoed": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_request[multi_admin_verify_request]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "approve_expiry_time": "string",
      "approve_time": "string",
      "approved_users": [
        "string"
      ],
      "comment": "string",
      "create_time": "string",
      "execute_on_approval": true,
      "execution_expiry_time": "string",
      "index": 0,
      "operation": "string",
```

```
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "pending_approvers": 0,
      "permitted_users": [
        "string"
      ],
      "potential_approvers": [
        "string"
      ],
      "query": "string",
      "required_approvers": 0,
      "state": "string",
      "user_requested": "string",
      "user_vetoed": "string"
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 262148
| The specified command is not recognized.

| 262279
| Invalid field set.

| 262304
| Too many requests. Delete one before creating another.

| 262305
| Can't approve non-pending request.

| 262306
| Can't veto an expired request.

| 262308
| The specified command is not supported by this feature.

| 262309
| The feature must be enabled first.

| 262311
| Value must be greater than zero.

| 262312
| Number of required approvers must be less than the total number of
unique approvers in the approval-groups.

| 262313
| Number of unique approvers in the approval-groups must be greater than
the number of required approvers.

| 262326
| Failed to parse query.

| 262327
| Failed to crate the request.

| 262328
| There is no matching rule for this request.

| 262330
| Cannot approve/veto a request multiple times.
```

```
| 262334
| The parameter specified in the command is not supported.

| 262337
| Cannot approve/veto the user's own request.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the request. This can identify the cluster or an SVM.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#multi_admin_verify_request]
[.api-collapsible-fifth-title]
multi_admin_verify_request

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.


|comment
|string
a|Optional user-provided comment that is sent to the approval-group email
indicating why the request was made.
```

```
|create_time
|string
a|

|execution_expiry_time
|string
a|

|index
|integer
a|Unique index that represents a request.


|operation
|string
a|The command to execute.


|owner
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|permitted_users
|array[string]
a|List of users that can execute the operation once approved. If not set,
any authorized user can perform the operation.


|potential_approvers
|array[string]
a|The users that are able to approve the request.


|query
|string
a|Identifies the specific entry upon which the user wants to operate.


|required_approvers
|integer
```

```
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
required.


|user_requested
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|user_vetoed
|string
a|The user that vetoed the request.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage a multi-admin-verify approval request
```

```
:leveloffset: +1


[[IDc2d93924bbb2f916f88d476fea3407d8]]
= Manage a multi-admin-verify approval request



== Overview

These APIs provide information about a specific multi-admin verification
request.
If you need to execute a command that is protected by a multi-admin rule,
you must first submit a request to be allowed to execute the command.
The request must then be approved by the designated approvers according to
the rule associated with the command.
 +

'''

== Examples

=== Retrieving a multi-admin-verify request

Retrieves information about a specific multi-admin verification request.
 +

'''

----

# The API:
/api/security/multi-admin-verify/requests/{index}

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-
verify/requests/1"

# The response:
{
"index": 1,
"operation": "security multi-admin-verify modify",
"query": "",
"state": "expired",
"required_approvers": 1,
"pending_approvers": 1,
"execute_on_approval": false,
```

```
"permitted_users": [
  "example_user"
],
"user_requested": "admin",
"owner": {
  "uuid": "c1483186-6e73-11ec-bc92-005056a7ad04",
  "name": "cluster1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/c1483186-6e73-11ec-bc92-005056a7ad04"
    }
  }
},
"create_time": "2022-01-05T20:07:09-05:00",
"approve_expiry_time": "2022-01-05T21:07:09-05:00",
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/requests/1"
  }
}
}
----

'''

=== Updating a multi-admin-verify request

Updates a specific multi-admin-verify request
 +

'''

----

# The API:
/api/security/multi-admin-verify/requests/{index}

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/multi-admin-
verify/requests/1" -d '{"state": "approved", "execute_on_approval":
false}'
----

'''
```

```
[[ID73293cdaf14b5ecf1d152950dae014bf]]
= Delete a multi-admin-verify request

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/multi-admin-verify/requests/{index}`#

*Introduced In:* 9.11

Deletes a multi-admin-verify request.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|index
|string
|path
|True
a|
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262336
```

```
| User not authorized to delete the request.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID9741e2c74f90ea00255e1cf9e080ed58]]
= Retrieve a multi-admin-verify request

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/requests/{index}`#

*Introduced In:* 9.11

Retrieves a multi-admin-verify request.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|index
|string
|path
|True
a|

|fields
|array[string]
|query
|False
```

```
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.


|comment
|string
a|Optional user-provided comment that is sent to the approval-group email
indicating why the request was made.


|create_time
|string
a|

|execute_on_approval
|boolean
a|Specifies that the operation is executed automatically on final
approval.


|execution_expiry_time
|string
a|
```

```
|index
|integer
a|Unique index that represents a request.


|operation
|string
a|The command to execute.


|owner
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|permitted_users
|array[string]
a|List of users that can execute the operation once approved. If not set,
any authorized user can perform the operation.


|potential_approvers
|array[string]
a|The users that are able to approve the request.


|query
|string
a|Identifies the specific entry upon which the user wants to operate.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
```

```
required.


|user_requested
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|user_vetoed
|string
a|The user that vetoed the request.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approve_expiry_time": "string",
  "approve_time": "string",
  "approved_users": [
    "string"
  ],
  "comment": "string",
  "create_time": "string",
  "execution_expiry_time": "string",
  "index": 0,
  "operation": "string",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "pending_approvers": 0,
  "permitted_users": [
```

```
    "string"
  ],
  "potential_approvers": [
    "string"
  ],
  "query": "string",
  "required_approvers": 0,
  "state": "string",
  "user_requested": "string",
  "user_vetoed": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
```

```
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
```

```
owner

The owner of the request. This can identify the cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDeb5b8cdebf2862efcf7e270ee451dfdb]]
= Update a multi-admin-verify request

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

```
block]#`/security/multi-admin-verify/requests/{index}`#

*Introduced In:* 9.11

Updates a multi-admin-verify request.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|index
|string
|path
|True
a|
|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.
```

```
|create_time
|string
a|

|execute_on_approval
|boolean
a|Specifies that the operation is executed automatically on final
approval.


|execution_expiry_time
|string
a|

|index
|integer
a|Unique index that represents a request.


|owner
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|potential_approvers
|array[string]
a|The users that are able to approve the request.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
required.
```

```
|user_requested
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|user_vetoed
|string
a|The user that vetoed the request.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approve_expiry_time": "string",
  "approve_time": "string",
  "approved_users": [
    "string"
  ],
  "create_time": "string",
  "execution_expiry_time": "string",
  "index": 0,
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "pending_approvers": 0,
  "potential_approvers": [
    "string"
  ],
  "required_approvers": 0,
  "state": "string",
  "user_requested": "string",
  "user_vetoed": "string"
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262309
| The feature must be enabled first.

| 262329
| Invalid state. State can be modified to approved, vetoed, or executed.

| 262337
| User not eligible to approve the request.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the request. This can identify the cluster or an SVM.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#multi_admin_verify_request]
[.api-collapsible-fifth-title]
multi_admin_verify_request

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|approve_expiry_time
|string
a|

|approve_time
|string
a|

|approved_users
|array[string]
a|The users that have approved the request.


|create_time
|string
a|

|execute_on_approval
|boolean
a|Specifies that the operation is executed automatically on final
approval.


|execution_expiry_time
|string
a|

|index
|integer
a|Unique index that represents a request.


|owner
|link:#owner[owner]
a|The owner of the request. This can identify the cluster or an SVM.


|pending_approvers
|integer
a|The number of approvers remaining that are required to approve.


|potential_approvers
|array[string]
a|The users that are able to approve the request.
```

```
|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|state
|string
a|The state of the request. PATCH supports approved and vetoed. The state
only changes after setting to approved once no more approvers are
required.


|user_requested
|string
a|The user that created the request. Automatically set by ONTAP.

* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|user_vetoed
|string
a|The user that vetoed the request.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1
```

```
= Manage multi-admin-verify rules

:leveloffset: +1


[[IDbac700d8db17f28482e320911c97101f]]
= Manage multi-admin-verify rules



== Overview

Rules define the ONTAP commands (operations) that should be protected by
multi-admin approval.
While the feature is turned on, any ONTAP operation that is defined with a
rule will be enforced with multi-admin approval to execute the command
(operation).
 +

'''

== Examples

=== Creating a multi-admin-verify rule

Creates a rule for the specified ONTAP operation.
 +

'''

----

# The API:
/api/security/multi-admin-verify/rules

# The call:
curl -X POST "https://<mgmt-ip>/api/security/multi-admin-
verify/rules?return_records=true" -H "accept: application/hal+json" -d
'{"operation": "volume delete", "query": "-vserver vs0",
"required_approvers": 1}'

# The response:
{
"num_records": 1,
"records": [
  {
    "owner": {
```

```
      "uuid": "c109634f-7011-11ec-a23d-005056a78fd5",
      "_links": {
        "self": {
          "href": "/api/svm/svms/c109634f-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "volume delete",
    "auto_request_create": true,
    "query": "-vserver vs0",
    "required_approvers": 1,
    "create_time": "2022-01-07T22:14:03-05:00",
    "system_defined": false,
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/c109634f-7011-
11ec-a23d-005056a78fd5/volume%20delete"
      }
    }
  }
]
}
----

'''

=== Retrieving multi-admin-verify rules

Displays information about multi admin verification rules.
 +

'''

----

# The API:
/api/security/multi-admin-verify/rules

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-verify/rules"

# The response:
{
"records": [
  {
    "owner": {
```

```
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security login password",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20login%20password"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security login unlock",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20login%20unlock"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify approval-group create",
    "_links": {
```

```
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20approval-
group%20create"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify approval-group delete",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20approval-
group%20delete"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify approval-group modify",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20approval-
group%20modify"
      }
    }
  },
  {
```

```
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify approval-group replace",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20approval-
group%20replace"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify modify",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20modify"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
```

```
    "operation": "security multi-admin-verify rule create",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20rule%20create"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify rule delete",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20rule%20delete"
      }
    }
  },
  {
    "owner": {
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "security multi-admin-verify rule modify",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/security%20multi-admin-verify%20rule%20modify"
      }
    }
  },
  {
    "owner": {
```

```
      "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
        }
      }
    },
    "operation": "volume delete",
    "_links": {
      "self": {
        "href": "/api/security/multi-admin-verify/rules/52b75787-7011-
11ec-a23d-005056a78fd5/volume%20delete"
      }
    }
  }
],
"num_records": 11,
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/rules"
  }
}
}
----

'''
```


```
[[IDf4a200a5a3936d3a7a4ee080203589f4]]
= Retrieve multi-admin-verify rules

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/rules`#

*Introduced In:* 9.11

Retrieves multi-admin-verify rules.


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|approval_expiry
|string
|query
|False
a|Filter by approval_expiry


|operation
|string
|query
|False
a|Filter by operation


|required_approvers
|integer
|query
|False
a|Filter by required_approvers


|query
|string
|query
|False
a|Filter by query


|execution_expiry
|string
|query
|False
a|Filter by execution_expiry


|system_defined
|boolean
|query
|False
a|Filter by system_defined
```

```
|auto_request_create
|boolean
|query
|False
a|Filter by auto_request_create


|approval_groups.name
|string
|query
|False
a|Filter by approval_groups.name


|create_time
|string
|query
|False
a|Filter by create_time


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid


|owner.name
|string
|query
|False
a|Filter by owner.name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
```

```
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_rule[multi_admin_verify_rule]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "approval_expiry": "string",
      "approval_groups": [
        {
          "name": "string"
        }
      ],
      "create_time": "string",
      "execution_expiry": "string",
      "operation": "string",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "query": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===
```

```
[#approval_groups]
[.api-collapsible-fifth-title]
approval_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the approval group.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the rule. The only valid owner is currently the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#multi_admin_verify_rule]
[.api-collapsible-fifth-title]
multi_admin_verify_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
operation where there is no matching pending request.
```

```
* Default value:
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|operation
|string
a|Command that requires one or more approvals.


|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.


|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
a|Specifies whether the rule is system-defined or user-defined.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID7af9938e299858679ec4c55bdcec2e05]]
= Create a multi-admin-verify rule

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/multi-admin-verify/rules`#

*Introduced In:* 9.11

Creates a multi-admin-verify rule.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
operation where there is no matching pending request.

* Default value: 1
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|operation
|string
a|Command that requires one or more approvals.


|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.
```

```
|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
a|Specifies whether the rule is system-defined or user-defined.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approval_expiry": "string",
  "approval_groups": [
    {
      "name": "string"
    }
  ],
  "create_time": "string",
  "execution_expiry": "string",
  "operation": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "query": "string"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#multi_admin_verify_rule[multi_admin_verify_rule]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {
      "approval_expiry": "string",
      "approval_groups": [
        {
          "name": "string"
        }
      ],
      "create_time": "string",
      "execution_expiry": "string",
      "operation": "string",
      "owner": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "query": "string"
    }
  ]
}
====
```

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262148
| The specified command is not recognized.

| 262308
| The specified command is not supported by this feature.

| 262309
| The feature must be enabled first.

| 262311
| Value must be greater than zero.

| 262312
| Number of required approvers must be less than the total number of
unique approvers in the approval-groups.

| 262313
| Number of unique approvers in the approval-groups must be greater than
```

```
the number of required approvers.

| 262314
| Some approval-groups were not found.

| 262316
| Value must be in the range two minutes to two weeks.

| 262326
| Failed to parse query.

| 262335
| The query string must be contained in either the "operation" or "query"
parameters but not in both.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#approval_groups]
[.api-collapsible-fifth-title]
approval_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the approval group.


|===


[#href]
```

```
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the rule. The only valid owner is currently the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#multi_admin_verify_rule]
[.api-collapsible-fifth-title]
multi_admin_verify_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
operation where there is no matching pending request.

* Default value: 1
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|operation
|string
a|Command that requires one or more approvals.
```

```
|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.


|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
a|Specifies whether the rule is system-defined or user-defined.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage a multi-admin-verify rule

:leveloffset: +1
```

```
[[ID68baa247fa883e7242c766c7bf41abed]]
= Manage a multi-admin-verify rule



== Overview

These APIs provide information about a specific multi-admin verification
rule.
Rules define the ONTAP commands (operations) that should be protected by
multi-admin approval.
While the feature is turned on, any ONTAP operation that is defined with a
rule will be enforced with multi-admin approval to execute the command
(operation).
 +

'''

== Examples

=== Retrieving a multi-admin-verify rule

Displays information about a specific multi admin verification rule.
 +

'''

----

# The API:
/api/security/multi-admin-verify/rules/{owner.uuid}/{operation}

# The call:
curl -X GET "https://<cluster-ip>/api/security/multi-admin-
verify/rules/52b75787-7011-11ec-a23d-005056a78fd5/volume+delete"

# The response:
{
"owner": {
  "uuid": "52b75787-7011-11ec-a23d-005056a78fd5",
  "name": "cluster1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/52b75787-7011-11ec-a23d-005056a78fd5"
    }
  }
},
```

```
"operation": "volume delete",
"auto_request_create": true,
"query": "-vserver vs0",
"required_approvers": 1,
"create_time": "2022-01-07T22:14:03-05:00",
"system_defined": false,
"_links": {
  "self": {
    "href": "/api/security/multi-admin-verify/rules/52b75787-7011-11ec-
a23d-005056a78fd5/volume+delete"
  }
}
}
----

'''

=== Updating a multi-admin-verify rule

Modifies the attributes of the rule.
 +

'''

----

# The API:
/api/security/multi-admin-verify/rules/{owner.uuid}/{operation}

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/multi-admin-
verify/rules/52b75787-7011-11ec-a23d-005056a78fd5/volume+delete" -d
'{"required_approvers": 1}'
----

'''

=== Deleting a multi-admin-verify rule

Deletes the specified approval group.
 +

'''

----

# The API:
```

```
/api/security/multi-admin-verify/rules/{owner.uuid}/{operation}

# The call:
curl -X DELETE "https://<cluster-ip>/api/security/multi-admin-
verify/rules/52b75787-7011-11ec-a23d-005056a78fd5/volume+delete"
----

'''





[[ID83d9eacdab4b3d23e49064908a05b0ba]]
= Delete a multi-admin-verify rule

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/multi-admin-verify/rules/{owner.uuid}/{operation}`#

*Introduced In:* 9.11

Deletes a multi-admin-verify rule.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|operation
|string
|path
|True
a|
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262310
| System rules cannot be deleted or have their query modified.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID00e78500a8783bdff2aed836a4014ea4]]
= Retrieve a multi-admin-verify rule

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/multi-admin-verify/rules/{owner.uuid}/{operation}`#

*Introduced In:* 9.11

Retrieves a multi-admin-verify rule.


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|operation
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
```

```
that don't have approval groups.


|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
operation where there is no matching pending request.

* Default value: 1
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|operation
|string
a|Command that requires one or more approvals.


|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.


|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
```

```
a|Specifies whether the rule is system-defined or user-defined.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "approval_expiry": "string",
  "approval_groups": [
    {
      "name": "string"
    }
  ],
  "create_time": "string",
  "execution_expiry": "string",
  "operation": "string",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "query": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#approval_groups]
[.api-collapsible-fifth-title]
approval_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the approval group.


|===
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the rule. The only valid owner is currently the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID911e16e2035bcdd1341ce35f32b6a6ee]]
= Update a multi-admin-verify rule

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/multi-admin-verify/rules/{owner.uuid}/{operation}`#

*Introduced In:* 9.11

Updates a multi-admin-verify rule.


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|

|operation
|string
|path
|True
a|
|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.


|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
```

```
operation where there is no matching pending request.

* Default value: 1
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.


|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
a|Specifies whether the rule is system-defined or user-defined.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "approval_expiry": "string",
  "approval_groups": [
    {
      "name": "string"
    }
  ],
  "create_time": "string",
  "execution_expiry": "string",
  "owner": {
    "name": "svm1"
  },
  "query": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262310
| System rules cannot be deleted or have their query modified.

| 262311
| Value must be greater than zero.

| 262312
| Number of required approvers must be less than the total number of
unique approvers in the approval-groups.

| 262313
| Number of unique approvers in the approval-groups must be greater than
the number of required approvers.

| 262316
| Value must be in the range two minutes to two weeks.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#approval_groups]
[.api-collapsible-fifth-title]
approval_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the approval group.


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

The owner of the rule. The only valid owner is currently the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a POST or PATCH
method.


|===


[#multi_admin_verify_rule]
[.api-collapsible-fifth-title]
multi_admin_verify_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|approval_expiry
|string
a|Time for requests to be approved, in ISO-8601 duration format. If not
set, the global setting is used.


|approval_groups
|array[link:#approval_groups[approval_groups]]
a|List of approval groups that are allowed to approve requests for rules
that don't have approval groups.
```

```
|auto_request_create
|boolean
a|When true, ONTAP automatically creates a request for any failed
operation where there is no matching pending request.

* Default value: 1
* Introduced in: 9.11
* x-nullable: true


|create_time
|string
a|

|execution_expiry
|string
a|Time for requests to be executed once approved, in ISO-8601 duration
format. If not set, the global setting is used.


|owner
|link:#owner[owner]
a|The owner of the rule. The only valid owner is currently the cluster.


|query
|string
a|When specified, this property limits the entries that require approvals
to those that match the specified query.


|required_approvers
|integer
a|The number of required approvers, excluding the user that made the
request.


|system_defined
|boolean
a|Specifies whether the rule is system-defined or user-defined.


|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage security roles

:leveloffset: +1


[[IDcdf8e2a5868a35e03c8717b9725bd426]]
= Manage security roles



== Overview

ONTAP supports Role Based Access Control (RBAC) wherein a user account
must be associated with a role and that role defines the privileges and
rights for that user account. A privilege defines the access level of the
API or command/command directory path. If a privilege tuple refers to a
command/command directory path, it can also be associated with an optional
query. The access level specifies the subset of operations a user account
can perform from the complete set of API methods {GET, POST, PATCH, and
DELETE} or command operations {create, delete, modify, and show}. The
optional query specifies the subset of objects that the role is allowed to
access. The query can be specified, only if the privilege tuple refers to
a command/command directory path. It is defined using one or more
parameters of the command/command directory path.

A role can comprise of multiple privilege tuples and each privilege tuple
consists of a REST API or command/command directory path, its access
level, and an optional query. For a given role, only one type of privilege
tuple can be defined. All privilege tuples for a role must contain REST
API paths or all privilege tuples for the role must contain
command/command directory paths. However, predefined/built-in roles (those
```

```
defined later) are an exception to this rule.

For example, "role1" might be a role that has a tuple
{"access":"all","path":"/api/network/ip"}, which means that a user account
with "role1" can perform GET, POST, PATCH, and DELETE requests on the
_api/network/ip_ API or derived APIs that have _api/network/ip_ as the
prefix.

In other examples, "role2" might be a role that has a tuple
{"access":"read_create_modify","path":"/api/storage/volumes"}, which means
that a user account with "role2" can perform GET, POST and PATCH (but not
DELETE) requests on the _api/storage/volumes_ API or derived APIs that
have _api/storage/volumes_ as the prefix.

"role3" might be a role that has a tuple
{"access":"read_create","path":"vserver nfs"}, which means that a user
account with "role3" can perform "show" and "create" operations on
_vserver nfs_ command or derived commands that have _vserver nfs_ as the
prefix. There is no query associated with "role3".

"role4" might be a role that has a tuple
{"access":"all","path":"snapmirror policy","query":"-policy !CustomPol*"},
which means that a user account with "role4" can perform "show",
"create","modify" and "delete" operations on _snapmirror policy_ command
or derived commands that have _snapmirror policy_ as the prefix. However,
a user is not authorized to perform the above set of operations on
SnapMirror policies starting with the name "CustomPol".

In cases where a role has tuples with multiple APIs having the same prefix
or multiple commands/command directories having the same prefix, the
highest match wins out. For example, if "role5" has the following tuples:
{"access":"readonly","path":"/api/cluster"} and
{"access":"all","path":"/api/cluster/schedules"}, then only a GET request
is allowed on APIs with _api/cluster_ as the prefix; while GET, POST,
PATCH and DELETE requests are possible on the _api/cluster/schedules_ API.
Similarly, if "role6" has the following tuples:
{"access":"readonly","path":"volume"} and
{"access":"read_create_delete","path":"volume snapshot"}, then only a
"show" operation is allowed on commands/command directories with _volume_
but not _volume snapshot_ as the prefix; while "show", "create" and
"delete" operations are possible on the _volume snapshot_ command
directory or any other command/command directory under _volume snapshot_.

=== Predefined (built-in) roles

Related REST APIs and related commands/command directories are used to
```

```
form predefined cluster-scoped and SVM-scoped roles, such as: "admin",
"backup", "readonly" for cluster and "vsadmin", "vsadmin-backup",
"vsadmin-protocol" for SVMs.

All of the above predefined roles can be retrieved by calling a GET
request on the _/api/security/roles_ API and can be assigned to user
accounts. See the examples for _api/security/accounts_.

A GET request on _/api/security/roles/{owner.uuid}/\{name}_ or
_/api/security/roles/{owner.uuid}/\{name}/privileges_, where "name" refers
to a predefined (built-in) role, returns privilege tuples containing REST
API paths along with privilege tuples containing command/command directory
paths.

These predefined roles cannot be modified or deleted.

=== Mapped roles

Before REST APIs, the RBAC roles (legacy roles) were defined to contain
the CLI commands and their access levels. Now, almost all REST APIs map to
one or more CLI commands. When a role is created using a POST request on
_/api/security/roles_, a mapped legacy role is created. This legacy role
has the same access level (as that of the REST API) for the mapped CLI
commands. However, if a legacy role with the same name already exists, the
POST operation fails and you need to choose a unique name for the role.
Legacy roles are also managed using the REST endpoint
_/api/security/roles_ and its derivatives. In CLI, legacy roles are
managed using the "security login role <create \| modify \| delete> -role
+++<rolename>+++" commands.+++</rolename>+++

Note that the mapped legacy role (for the REST API role created) cannot be
manipulated using either REST API or the CLI.

The reverse case is not true; the creation of a legacy role will not
create a mapped role with equivalent REST APIs.

=== API restrictions

A role can be a REST role or a legacy role but not both. A role cannot be
defined to have a mix of privilege tuples with REST API paths and
privilege tuples with command/command directory paths. However, predefined
(built-in) roles are an exception to this rule.
Numerous APIs are scoped for the cluster level only. This results in an
access error if assigned to an SVM-scoped role. For example,
_api/cluster/nodes_ does not work when added as a tuple entry for an SVM-
scoped role.
```

```
A number of APIs allowed for an SVM-scoped role might have restrictions on
the access level. For example, _/api/network/ethernet/ports_ cannot have
an access level of "all" for an SVM-scoped role; this results in an access
error when a POST or PATCH request is made.

Roles created with a REST API path prefix which is common to many APIs
might have restrictions based on the scope of the role; cluster or SVM.
For example, {"access":"all","path":"/api/security"} might be a tuple
entry for an SVM role. Any GET, POST, PATCH, or DELETE operation fails on
API _/api/security/accounts_ while the same on
_/api/security/login/messages_ succeeds. However, a role with exactly the
same tuple when created at the cluster-scope level allows the operations.

Numerous APIs have restrictions on the objects that can be operated on
based on the context of the SVM or cluster. For example, a POST request on
_/api/security/authentication/password_ API changes the password for a
user account. If executed in the context of an SVM (POST request on an SVM
interface), only the password of the user executing the POST can be
modified, and attempts to modify the password of any other user results in
an access error. However, if a POST request is performed by a cluster
administrator account, the password for any user account (cluster or SVM)
can be modified.

Resource-qualified endpoints are now supported. At present, the only
supported resource-qualified endpoints are the following&#58;

==== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

==== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_
```

```
&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

==== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

== Examples

=== Creating a cluster-scoped custom role of REST API tuples

Specify the role name and the tuples (of REST APIs and their access
levels) in the body of the POST request. The owner.uuid or owner.name are
not required to be specified for a cluster-scoped role.

----

# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d
'{"name":"cluster_role1", "privileges" :
[{"access":"readonly","path":"/api/cluster/jobs"},{"access":"all","path":"
/api/application/applications"},{"access":"readonly","path":"/api/applicat
ion/templates"}]}'
----

=== Creating a cluster-scoped custom role of command and/or command
directory tuples

Specify the role name and the tuples (of commands/command directories,
their access levels and associated optional queries) in the body of the
POST request. The owner.uuid or owner.name are not required to be
specified for a cluster-scoped role.

----

# The API:
```

```
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d
'{"name":"cluster_role2", "privileges" :
[{"access":"readonly","path":"volume
qtree","query":""},{"access":"all","path":"security
certificate"},{"access":"readonly","path":"snapmirror policy","query":"-
policy !CustomPol*"}]}'
----

=== Creating an SVM-scoped custom role of REST API tuples

For an SVM-scoped role, specify either owner.name or owner.uuid in the
request body along with other parameters for the role. These correspond to
the name or UUID of the SVM for which the role is being created and can be
obtained from the response body of the GET request performed on the
_/api/svm/svms_ API.

----

# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d '{"owner": {"uuid"
: "9f93e553-4b02-11e9-a3f9-005056bb7acd"},"name":"svm_role1", "privileges"
:
[{"access":"readonly","path":"/api/cluster/jobs"},{"access":"all","path":"
/api/application/applications"},{"access":"readonly","path":"/api/applicat
ion/templates"}]}'
----

=== Creating an SVM-scoped custom role of command and/or command directory
tuples

For an SVM-scoped role, specify either owner.name or owner.uuid in the
request body along with other parameters for the role. These correspond to
the name or UUID of the SVM for which the role is being created and can be
obtained from the response body of the GET request performed on the
_/api/svm/svms_ API.

----

# The API:
POST "/api/security/roles"
```

```
# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d '{"owner": {"uuid"
: "9f93e553-4b02-11e9-a3f9-005056bb7acd"},"name":"svm_role2", "privileges"
: [{"access":"readonly","path":"job schedule interval","query":"-days
>1"},{"access":"all","path":"application
snapshot"},{"access":"none","path":"volume move"}]}'
----

=== Creating a custom role with a resource-qualified endpoint

Specify the role name and the tuples (of REST APIs and their access
levels) in the body of the POST request. One or more of the tuples can now
contain a resource-qualified endpoint. At present, the only supported
resource-qualified endpoints are the _Snapshots_, _ONTAP S3_, and _File
System Analytics_ endpoints listed above in the _Overview_ section.

----

# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d
'{"name":"cluster_role", "privileges" :
[{"access":"readonly","path":"/api/cluster/jobs"},{"access":"all","path":"
/api/storage/volumes/4ae77149-7752-11eb-8d4e-
0050568ed6bd/snapshots"},{"access":"all","path":"/api/storage/volumes/6519
986e-7752-11eb-8d4e-
0050568ed6bd/snapshots"},{"access":"readonly","path":"/api/storage/volumes
/8823c869-9ea1-11ec-8771-005056bb1a7c/top-
metrics/users"},{"access":"readonly","path":"/api/application/templates"}]
}'
----

=== Creating a custom role with a private CLI endpoint

----

# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d
'{"name":"cluster_role3", "privileges" :
[{"access":"readonly","path":"/api/private/cli/cluster"}]}'
----
```

```
=== Retrieving the configured roles

All of the roles or a filtered list of roles (for example by name,
predefined, and so on) can be retrieved.

----

# The API:
GET "/api/security/roles?fields=%2A"

# The call to retrieve all the roles configured in the cluster:
curl -X GET "https://<mgmt-ip>/api/security/roles"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "admin",
    "privileges": [
      {
        "path": "/api",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/admin/privileges/%2Fapi"
          }
        }
      },
      {
        "path": "DEFAULT",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/admin/privileges/DEFAULT"
          }
```

```
        }
      }
    ],
    "builtin": true,
    "scope": "cluster",
    "_links": {
      "self": {
        "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/admin"
      }
    }
  },
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "customRole_rest",
    "privileges": [
      {
        "path": "/api/storage/volumes/738e3c9f-9897-41f2-be92-
a00945fd9bdb/snapshots",
        "access": "readonly",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_rest/privileges/%2Fapi%2Fstorage%2Fvolumes%2F738e3
c9f-9897-41f2-be92-a00945fd9bdb%2Fsnapshots"
          }
        }
      },
      {
        "path": "/api/storage/volumes/e621583b-f445-4713-ba9e-
a052d53c8a83/snapshots",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_rest/privileges/%2Fapi%2Fstorage%2Fvolumes%2Fe6215
83b-f445-4713-ba9e-a052d53c8a83%2Fsnapshots"
          }
```

```
        }
      }
      {
        "path": "/api/svm/svms/881764b5-9ea1-11ec-8771-005056bb1a7c/top-
metrics/directories",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/881764b5-9ea1-11ec-8771-
005056bb1a7c/customRole_rest/privileges/%2Fapi%2Fstorage%2Fsvm%2F881764b5-
9ea1-11ec-8771-005056bb1a7c%2Ftop-metrics%2Fdirectories"
          }
        }
      }
    ],
    "builtin": false,
    "scope": "cluster",
    "_links": {
      "self": {
        "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_rest"
      }
    }
  },
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "customRole_legacy",
    "privileges": [
      {
        "path": "volume",
        "access": "readonly",
        "query":"-is_svm_root false",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_legacy/privileges/volume"
          }
        }
```

```
      },
      {
        "path": "volume snapshot",
        "access": "all",
        "query": "-volume vol1&#124;vol2",
        "_links": {
          "self": {
            "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_legacy/privileges/volume%20snapshot"
          }
        }
      }
    ],
    "builtin": false,
    "scope": "cluster",
    "_links": {
      "self": {
        "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-
0050568e2e25/customRole_legacy"
      }
    }
  },
  {
    "owner": {
      "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "vsadmin",
    "privileges": [
      {
        "path": "/api/application/applications",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/%2Fapi%2Fapplication%2Fapplications"
          }
        }
      },
      {
        "path": "/api/application/templates",
```

```
        "access": "readonly",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/%2Fapi%2Fapplication%2Ftemplates"
          }
        }
      },
      {
        "path": "/api/cluster",
        "access": "readonly",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/%2Fapi%2Fcluster"
          }
        }
      },
      {
        "path": "/api/cluster/jobs",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/%2Fapi%2Fcluster%2Fjobs"
          }
        }
      },
      {
        "path": "/api/cluster/schedules",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/%2Fapi%2Fcluster%2Fschedules"
          }
        }
      },
      {
        "path": "DEFAULT",
        "access": "none",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/DEFAULT"
          }
```

```
        }
      },
      {
        "path": "application create",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/application%20create"
          }
        }
      },
      {
        "path": "application delete",
        "access": "all",
        "_links": {
          "self": {
            "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin/privileges/application%20delete"
          }
        }
      },
    ],
    "builtin": true,
    "scope": "svm",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/vsadmin"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/security/roles?fields=%2A"
  }
}
}
----

=== Using a scoped call to retrieve the configured roles

----
```

```
# Scoped call to retrieve all the roles for a particular SVM using
owner.uuid:
curl -X GET "https://<mgmt-ip>/api/security/roles/?owner.uuid=aaef7c38-
4bd3-11e9-b238-0050568e2e25"

# Scoped call to retrieve all the roles for a particular SVM using
owner.name:
curl -X GET "https://<mgmt-ip>/api/security/roles/?owner.name=svm1"

# Scoped call to retrieve the roles having vsadmin as the prefix in the
role name:
curl -X GET "https://<mgmt-ip>/api/security/roles/?name=vsadmin*"

# Scoped call to retrieve the predefined roles:
curl -X GET "https://<mgmt-ip>/api/security/roles/?builtin=true"

# Scoped call to retrieve the custom roles:
curl -X GET "https://<mgmt-ip>/api/security/roles/?builtin=false"
----




[[ID9a548e96ad23536f819351c6419f51f4]]
= Retrieve a list of roles configured in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/roles`#

*Introduced In:* 9.6

Retrieves a list of roles configured in the cluster.

== Related ONTAP commands

* `security login rest-role show`
* `security login role show`

== Learn more

* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.name
|string
|query
|False
a|Filter by owner.name

* Introduced in: 9.7


|owner.uuid
|string
|query
|False
a|Filter by owner.uuid

* Introduced in: 9.7


|name
|string
|query
|False
a|Filter by name

* Introduced in: 9.7


|builtin
|boolean
|query
|False
a|Filter by builtin

* Introduced in: 9.7


|scope
|string
```

```
|query
|False
a|Filter by scope

* Introduced in: 9.7


|privileges.path
|string
|query
|False
a|Filter by privileges.path

* Introduced in: 9.7


|privileges.access
|string
|query
|False
a|Filter by privileges.access

* Introduced in: 9.7


|privileges.query
|string
|query
|False
a|Filter by privileges.query

* Introduced in: 9.11


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|num_records
|integer
a|Number of records


|records
|array[link:#role[role]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "admin",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "privileges": [
```

```
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "access": "all",
          "path": "volume move start",
          "query": "-vserver vs1|vs2|vs3 -destination-aggregate
aggr1|aggr2"
        }
      ],
      "scope": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
```

```
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
```

```
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs

* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#role]
[.api-collapsible-fifth-title]
role

A named set of privileges that defines the rights an account has when it
is assigned the role.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|builtin
|boolean
a|Indicates if this is a built-in (pre-defined) role which cannot be
modified or deleted.


|name
|string
a|Role name


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the role.


|privileges
|array[link:#role_privilege[role_privilege]]
a|The list of privileges that this role has been granted.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
```

```
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDd7d1b4b2b39d1ba88b17dfa8f8c77732]]
= Create a new cluster-scoped or SVM-scoped role

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/roles`#

*Introduced In:* 9.6

Creates a new cluster-scoped role or an SVM-scoped role. For an SVM-scoped
role, specify either the SVM name as the owner.name or SVM UUID as the
owner.uuid in the request body along with other parameters for the role.
The owner.uuid or owner.name are not required to be specified for a
cluster-scoped role.

== Required parameters

* `name` - Name of the role to be created.
* `privileges` - Array of privilege tuples. Each tuple consists of a REST
API or command/command directory path and its desired access level. If the
tuple refers to a command/command directory path, it could optionally
contain a query.

== Optional parameters

* `owner.name` or `owner.uuid`  - Name or UUID of the SVM for an SVM-
scoped role.

== Related ONTAP commands

* `security login rest-role create`
* `security login role create`

== Learn more

* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|builtin
|boolean
a|Indicates if this is a built-in (pre-defined) role which cannot be
modified or deleted.


|name
|string
a|Role name


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the role.


|privileges
|array[link:#role_privilege[role_privilege]]
a|The list of privileges that this role has been granted.


|scope
|string
```

```
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "admin",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "privileges": [
    {
      "access": "all",
      "path": "volume move start",
      "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
    }
  ],
  "scope": "string"
}
====

== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1263347
| Cannot modify pre-defined roles

| 2621462
| The supplied SVM does not exist.

| 5636129
| Role with given name has not been defined.

| 5636143
| Vserver admin cannot use the API with this access level.

| 5636144
| Invalid value specified for access level.

| 5636168
| This role is mapped to a REST role and can only be modified by updating
the REST role.
```

```
| 5636169
| Invalid character in URI.

| 5636170
| URI does not exist.

| 5636171
| Role already exists in legacy role table.

| 5636184
| Expanded REST roles for granular resource control feature is currently
disabled.

| 5636185
| The specified UUID was not found.

| 5636186
| Expanded REST roles for granular resource control requires an effective
cluster version of 9.10.1 or later.

| 5636191
| The "path" parameter in a "privileges" tuple can contain only API
endpoint entries or only command and command directory entries.

| 5636192
| The query parameter cannot be specified for the privileges tuple with
API endpoint entries.

| 5636200
| The specified value of the access parameter is invalid, if a command or
command directory is specified in the path parameter.

| 5636210
| User creation failed because LDAP is not configured for the SVM or the
LDAP connection is not secure.

| 5636262
| Cannot create a role with the specified role name because it is reserved
by the system.

| 13434890
| Vserver-ID failed for Vserver roles.

| 13434891
| UUID lookup failed for Vserver roles.
```

```
| 13434892
| Roles is a required field.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the role.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs

* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
```

```
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#role]
[.api-collapsible-fifth-title]
role

A named set of privileges that defines the rights an account has when it
is assigned the role.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|builtin
|boolean
a|Indicates if this is a built-in (pre-defined) role which cannot be
modified or deleted.


|name
|string
a|Role name


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the role.


|privileges
|array[link:#role_privilege[role_privilege]]
a|The list of privileges that this role has been granted.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
```

```
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= View or delete a role

:leveloffset: +1


[[ID9d57a42a7a0fdf0bc2cde2715b254cb5]]
= View or delete a role



== Overview

This API is used to retrieve or delete a role. The role can be SVM-scoped
or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID
corresponds to the UUID of the SVM for which the role has been created and
can be obtained from the response body of a GET call performed on one of
the following APIs:
_/api/security/roles_ for all roles
_/api/security/roles/?scope=svm_ for SVM-scoped roles
_/api/security/roles/?owner.name=+++<svm-name>+++</i> for roles in a
specific SVM This API response contains the complete URI for each role
that can be used for retrieving or deleting a role.+++</svm-name>+++_

NOTE: The pre-defined roles can be retrieved but cannot be deleted.

== Examples

=== Retrieving the role configuration for a REST role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/secure_role"

# The response:
{
```

```
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
    }
  }
},
"name": "secure_role",
"privileges": [
  {
    "path": "/api/security",
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/secure_role/privileges/%2Fapi%2Fsecurity"
      }
    }
  },
  {
    "path": "/api/storage/volumes/651f7fdf-7752-11eb-8d4e-
0050568ed6bd/snapshots",
    "access": "readonly",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/secure_role/privileges/%2Fapi%2Fstorage%2Fvolumes%2F651f7fdf-
7752-11eb-8d4e-0050568ed6bd%2Fsnapshots"
      }
    }
  },
  {
    "path": "/api/storage/volumes/6dfeb406-9a16-11ec-819e-
005056bb1a7c/top-metrics/clients",
    "access": "readonly",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/secure_role/privileges/%2Fapi%2Fstorage%2Fvolumes%2F6dfeb406-
9a16-11ec-819e-005056bb1a7c%2Ftop-metrics%2Fclients"
      }
    }
  }
],
```

```
"builtin": false,
"scope": "svm",
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/secure_role"
  }
}
}
----

=== Retrieving the role configuration for a custom legacy role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/finVolNoDel"

# The response:
{
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
    }
  }
},
"name": "finVolNoDel",
"privileges": [
  {
    "path": "DEFAULT",
    "access": "none",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/finVolNoDel/privileges/DEFAULT"
      }
    }
  },
  {
```

```
    "path": "volume",
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/finVolNoDel/privileges/volume"
      }
    }
  },
  {
    "path": "volume delete",
    "access": "none",
    "query": "-volume vol_fin*",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/finVolNoDel/privileges/volume%20delete"
      }
    }
  }
],
"builtin": false,
"scope": "svm",
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/finVolNoDel"
  }
}
}
----

=== Deleting a custom role

----

# The API:
DELETE "/api/security/roles/{owner.uuid}/{name}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1"
----
```

```
[[ID53ab07793381b466e731eb0eecf5fc51]]
= Delete a role

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}`#

*Introduced In:* 9.6

Deletes the specified role.

== Required parameters

* `name` - Name of the role to be deleted.
* `owner.uuid` - UUID of the SVM housing the role.

== Related ONTAP commands

* `security login rest-role delete`
* `security login role delete`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_endpoint_overview.html[
DOC /security/roles/{owner.uuid}/\{name}]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID
```

```
|name
|string
|path
|True
a|Role name to be deleted.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1263347
| Cannot modify pre-defined roles.

| 5636169
| Specified URI path is invalid or not supported. Resource-qualified
endpoints are not supported.

| 5636170
| URI does not exist.

| 5636172
| User accounts detected with this role assigned. Update or delete those
accounts before deleting this role.

| 5636173
| Features require an effective cluster version of 9.6 or later.

| 5636184
| Expanded REST roles for granular resource control feature is currently
disabled.

| 5636185
| The specified UUID was not found.
```

```
| 5636186
| Expanded REST roles for granular resource control requires an effective
cluster version of 9.10.1 or later.

| 5636245
| Failed to delete the role because it is part of the group role-mapping
configuration.

| 5636246
| Failed to delete the rest role because it is part of the group role-
mapping configuration.

| 5636247
| Failed to delete the rest role because it is part of the external-role-
mapping configuration.

| 5636249
| Failed to delete the role because it is part of the external-role-
mapping configuration.

| 5636259
| The specified child AIDE object was not found within the specified
parent AIDE object.

| 5636261
| The specified grandchild AIDE object was not found within the specified
child AIDE object that belongs to the specified parent AIDE
object.</personalities>

| 13434890
| Vserver-ID failed for Vserver roles.

| 13434893
| The SVM does not exist.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID3d541af1d19568c74ac49104f8b64d16]]
= Retrieve the details of a role

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}`#

*Introduced In:* 9.6

Retrieves the details of the specified role.

== Related ONTAP commands

* `security login rest-role show`
* `security login role show`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_endpoint_overview.html[
DOC /security/roles/{owner.uuid}/\{name}]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID


|name
|string
```

```
|path
|True
a|Role name


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|builtin
|boolean
a|Indicates if this is a built-in (pre-defined) role which cannot be
modified or deleted.


|name
|string
a|Role name


|owner
|link:#owner[owner]
a|Owner name and UUID that uniquely identifies the role.


|privileges
|array[link:#role_privilege[role_privilege]]
a|The list of privileges that this role has been granted.
```

```
|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "admin",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "privileges": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access": "all",
      "path": "volume move start",
      "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
    }
  ],
  "scope": "string"
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#owner]
[.api-collapsible-fifth-title]
owner

Owner name and UUID that uniquely identifies the role.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs

* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
```

```
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage role privilege details

:leveloffset: +1


[[IDbbfbeaa0d71c24de471af3c4ffffad99]]
= Manage role privilege details



== Overview

This API is used to configure the role privileges (tuples of REST URI
paths or command/command directory paths, their access levels and optional
queries, where the tuples refer to command/command directory paths). It
also retrieves all of the privilege tuples for a role and can add a tuple
to an existing role or creates a new role with the provided tuple. The
"path" attribute can be a resource-qualified endpoint. Currently, the only
supported resource-qualified endpoints are the following&#58;

=== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

=== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_
```

```
&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

==== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

The role can be SVM-scoped or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID
corresponds to the UUID of the SVM for which the role has been created and
can be obtained from the response body of a GET request performed on one
of the following APIs&#58;

* _/api/security/roles_ for all the roles
* _/api/security/roles/?scope=svm_ for SVM-scoped roles
* _/api/security/roles/?owner.name=+++<svm-name>+++</i> for roles in a
specific SVM This API response contains the complete URI for each role and
can be used after suffixing it with _"privileges"_.+++</svm-name>+++_

NOTE: The pre-defined roles can be retrieved but cannot be updated.

== Examples

=== Adding a privilege tuple for a REST URI/endpoint to an existing custom
role

----

# The API:
POST "/security/roles/{owner.uuid}/{name}/privileges"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
```

```
b238-0050568e2e25/svm_role1/privileges" -d
'{"access":"readonly","path":"/api/protocols"}'
----

=== Adding a privilege tuple for a command or command directory to an
existing custom role

----

# The API:
POST "/security/roles/{owner.uuid}/{name}/privileges"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges" -d
'{"access":"all","path":"statistics volume show","query":"-vserver
vs1&#124;vs2 -aggregate aggr1&#124;aggr2"}'
----

=== Retrieving all the privilege tuples for a REST role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges"

# The response:
{
"records": [
  {
    "path": "/api/application",
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fapplication"
      }
    }
  },
  {
    "path": "/api/protocols",
    "access": "readonly",
    "_links": {
```

```
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
      }
    }
  },
  {
    "path": "/api/storage/volumes/1385d680-74fc-4adb-a348-
9a740e83702a/snapshots",
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fvolumes%2F1385d680-
74fc-4adb-a348-9a740e83702a%2Fsnapshots"
      }
    }
  },
  {
    "path": "/api/storage/volumes/*/top-metrics/users",
    "access": "read_create_modify",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fvolumes%2F%2A%2Ftop-
metrics%2Fusers"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges"
  }
}
}
```


---

=== Retrieving all the privilege tuples for a custom legacy role

---

```
# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges"
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges"

# The response:
{
"records": [
  {
    "path": "network interface",
    "access": "readonly",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/network%20interface"
      }
    }
  },
  {
    "path": "security",
    "access": "readonly",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/security"
      }
    }
  },
  {
    "path": "security certificate"
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/security%20certificate"
      }
    }
  },
  {
    "path": "security password"
    "access": "all",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/security%20password"
      }
```

```
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges"
  }
}
}
----
```


```
[[ID434b6768113aee9cff55cc99a608d4b9]]
= Retrieve privilege details of the specified role

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}/privileges`#

*Introduced In:* 9.6

Retrieves privilege details of the specified role.

== Related ONTAP commands

* `security login rest-role show`
* `security login role show`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_privileges_endpoint_ove
rview.html[DOC /security/roles/{owner.uuid}/\{name}/privileges]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID


|name
|string
|path
|True
a|Role name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
```

```
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#role_privilege[role_privilege]]
a|

|===


.Example response
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access": "all",
      "path": "volume move start",
      "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs
```

```
* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
```

```
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDe8954d8d4eecf7b94b080c2b24b55378]]
= Add a privilege tuple to an existing role

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}/privileges`#

*Introduced In:* 9.6

Adds a privilege tuple (of REST URI or command/command directory path, its
access level and an optional query, if the "path" refers to a
command/command directory path) to an existing role or creates a new role
with the provided tuple.

== Required parameters

* `owner.uuid` - UUID of the SVM that houses this role.
* `name` - Name of the role to be updated.
* `path` - REST URI path (example: _/api/storage/volumes_) or
command/command directory path (example: _snaplock compliance-clock_). Can
be a resource-qualified endpoint (example: _/api/storage/volumes/43256a71-
```

```
be02-474d-a2a9-9642e12a6a2c/snapshots_). Currently, resource-qualified
endpoints are limited to the following&#58;

== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

* `access` - Desired access level for the REST URI path or command/command
directory.

== Related ONTAP commands

* `security login rest-role create`
* `security login role create`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_privileges_endpoint_ove
```

```
rview.html[DOC /security/roles/{owner.uuid}/\{name}/privileges]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID


|name
|string
|path
|True
a|Role name


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access": "all",
  "path": "volume move start",
  "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
}
====

== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1263347
| Cannot modify pre-defined roles.

| 5636129
| A role with given name has not been defined.

| 5636143
| A Vserver admin cannot use the API with this access level.

| 5636144
| The value specified for the access level is not valid.

| 5636168
| This role is mapped to a REST role and can only be modified by updating
the REST role.

| 5636262
| Cannot create a role with the specified role name because it is reserved
```

```
by the system.

| 13434890
| Vserver-ID failed for Vserver roles.

| 13434891
| UUID LookUp failed for Vserver roles.

| 13434892
| Roles is a required field.

| 13434893
| The SVM does not exist.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs

* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
```

```
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage role privilege path

:leveloffset: +1


[[IDa2f39e7426fd6d26a905a4f08ef365f5]]
= Manage role privilege path
```

```
== Overview

A role can comprise of multiple tuples and each tuple consists of a REST
API path or command/command directory path and its access level. If the
tuple refers to a command/command directory path, it may optionally be
associated with a query. These APIs can be used to retrieve or modify the
associated access level and optional query. They can also be used to
delete one of the constituent REST API paths or command/command directory
paths within a role. The REST API path can be a resource-qualified
endpoint. Currently, the only supported resource-qualified endpoints are
the following&#58;

=== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

=== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash;  _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

==== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
```

```
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

The role can be SVM-scoped or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID
corresponds to the UUID of the SVM for which the role has been created and
can be obtained from the response body of a GET request performed on one
of the following APIs&#58;
_/api/security/roles_ for all roles
_/api/security/roles/?scope=svm_ for SVM-scoped roles
_/api/security/roles/?owner.name=+++<svm-name>+++</i> for roles in a
specific SVM This API response contains the complete URI for each tuple of
the role and can be used for GET, PATCH, or DELETE operations.+++</svm-
name>+++_

NOTE: The access level for paths in pre-defined roles cannot be updated.

== Examples

=== Updating the access level for a REST API path in the privilege tuple
of an existing role

----

# The API:
PATCH "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols" -d
'{"access":"all"}'
----

=== Updating the access level for a command/command directory path in the
privilege tuple of an existing role

----

# The API:
PATCH "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/netp%20port" -d
'{"access":"readonly","query":"-type if-group&#124;vlan"}'
```

```
----

=== Updating the access level for a resource-qualified endpoint in the
privilege tuple of an existing role

----

# The API:
PATCH "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fvolumes%2F742ef001-
24f0-4d5a-9ec1-2fdaadb282f4%2Ffiles" -d '{"access":"readonly"}'
----

=== Retrieving the access level for a REST API path in the privilege tuple
of an existing role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"

# The response:
{
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25"
},
"name": "svm_role1",
"path": "/api/protocols",
"access": "all",
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
  }
}
}
----

=== Retrieving the access level for a command/command directory path in
```

```
the privilege tuple of an existing role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/net%20port"

# The response:
{
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25"
},
"name": "svm_role1",
"path": "net port",
"query":"-type if-group&#124;vlan",
"access": "readonly",
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/net%20port"
  }
}
}
----

=== Retrieving the access level for a resource-qualified endpoint in the
privilege tuple of an existing role

----

# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fvolumes%2Fd0f3b91a-
4ce7-4de4-afb9-7eda668659dd%2F%2Fsnapshots"

# The response:
{
"owner": {
  "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25"
```

```
},
"name": "svm_role1",
"path": "/api/storage/volumes/d0f3b91a-4ce7-4de4-afb9-
7eda668659dd/snapshots",
"access": "all",
"_links": {
  "self": {
    "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-
0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fvolumes%2Fd0f3b91a-
4ce7-4de4-afb9-7eda668659dd%2Fsnapshots"
  }
}
}
----

=== Deleting a privilege tuple, containing a REST API path, from an
existing role

----

# The API:
DELETE "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
----

=== Deleting a privilege tuple, containing a command/command directory
path, from an existing role

----

# The API:
DELETE "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/net%20port"
----

=== Deleting a privilege tuple, containing a resource-qualified endpoint,
from an existing role

----

# The API:
```

```
DELETE "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fstorage%2Fsvm%2F6e000659-
9a16-11ec-819e-005056bb1a7c%2Ftop-metrics%2Ffiles"
----




[[IDc1b216eda263e94aaa87424e0cddf472]]
= Delete a privilege tuple from the role

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}/privileges/{path}`#

*Introduced In:* 9.6

Deletes a privilege tuple (of REST URI or command/command directory path,
its access level and an optional query) from the role. The REST URI can be
a resource-qualified endpoint. Currently, the only supported resource-
qualified endpoints are the following&#58;

== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_
```

```
&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

== Required parameters

* `owner.uuid` - UUID of the SVM which houses this role.
* `name` - Name of the role to be updated.
* `path` - Constituent REST API path or command/command directory path to
be deleted from this role. Can be a resource-qualified endpoint (example:
_/api/svm/svms/43256a71-be02-474d-a2a9-9642e12a6a2c/top-metrics/users_).
Currently, resource-qualified endpoints are limited to the _Snapshots_,
_File System Analytics_ and _ONTAP S3_ endpoints listed above in the
description.

== Related ONTAP commands

* `security login rest-role delete`
* `security login role delete`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_privileges_path_endpoin
t_overview.html[DOC
/security/roles/{owner.uuid}/\{name}/privileges/\{path}]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID


|name
|string
|path
|True
a|Role name


|path
|string
|path
|True
a|REST API path or command/command directory path

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1263347
| Cannot modify pre-defined roles.

| 5636168
| This role is mapped to a REST role and can only be modified by updating
the REST role.
```

```
| 13434890
| Vserver-ID failed for Vserver roles.

| 13434893
| The SVM does not exist.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID247f310cff406af9691d20318a393708]]
= Retrieve the access level for a REST API path or command/command
directory path for a role

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}/privileges/{path}`#

*Introduced In:* 9.6

Retrieves the access level for a REST API path or command/command
directory path for the specified role. Optionally retrieves the query, if
'path' refers to a command/command directory path. The REST API path can
be a resource-qualified endpoint. Currently, the only supported resource-
qualified endpoints are the following&#58;

== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
```

```
&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

== Related ONTAP commands

* `security login rest-role show`
* `security login role show`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_privileges_path_endpoin
t_overview.html[DOC
/security/roles/{owner.uuid}/\{name}/privileges/\{path}]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID


|name
|string
|path
|True
a|Role name


|path
|string
|path
|True
a|REST API path or command/command directory path


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|path
|string
a|Either of REST URI/endpoint OR command/command directory path.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access": "all",
  "path": "volume move start",
  "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6a95b026eb13b5481e8909a6e20f9e8a]]
= Update the access level for a REST API path or command/command directory
path

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/roles/{owner.uuid}/{name}/privileges/{path}`#
```

```
*Introduced In:* 9.6

Updates the access level for a REST API path or command/command directory
path. Optionally updates the query, if 'path' refers to a command/command
directory path. The REST API path can be a resource-qualified endpoint.
Currently, the only supported resource-qualified endpoints are the
following&#58;

== Snapshots APIs

&ndash; _/api/storage/volumes/{volume.uuid}/snapshots_

== File System Analytics APIs

&ndash; _/api/storage/volumes/{volume.uuid}/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/clients_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/directories_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/files_

&ndash; _/api/storage/volumes/{volume.uuid}/top-metrics/users_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/clients_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/directories_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/files_

&ndash; _/api/svm/svms/{svm.uuid}/top-metrics/users_

== ONTAP S3 APIs

&ndash; _/api/protocols/s3/services/{svm.uuid}/users_

In the APIs above, and in the context of REST roles, the wildcard
character &#42; can be used in place of _{volume.uuid}_ or _{svm.uuid}_ to
represent _all_ volumes or _all_ SVMs, depending on whether the REST
endpoint references volumes or SVMs. The _{volume.uuid}_ corresponds to
the _-instance-uuid_ field in the output of the "volume show" command at
the diagnostic privilege level. It can also be retrieved through the REST
endpoint _/api/storage/volumes_.

== Required parameters
```

```
* `owner.uuid` - UUID of the SVM that houses this role.
* `name` - Name of the role to be updated.
* `path` - Constituent REST API path or command/command directory path,
whose access level and/or query are/is to be updated. Can be a resource-
qualified endpoint (example: _/api/storage/volumes/43256a71-be02-474d-
a2a9-9642e12a6a2c/snapshots_). Currently, resource-qualified endpoints are
limited to the _Snapshots_, _File System Analytics_ and _ONTAP S3_
endpoints listed above in the description.

== Optional parameters

* `access` - Access level for the path.
* `query` - Optional query, if the path refers to a command/command
directory path.

== Related ONTAP commands

* `security login rest-role modify`
* `security login role modify`

== Learn more

*
xref:{relative_path}security_roles_owner.uuid_name_privileges_path_endpoin
t_overview.html[DOC
/security/roles/{owner.uuid}/\{name}/privileges/\{path}]
* xref:{relative_path}security_roles_endpoint_overview.html[DOC
/security/roles]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Role owner UUID
```

```
|name
|string
|path
|True
a|Role name


|path
|string
|path
|True
a|REST API path or command/command directory path

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access": "all",
  "query": "-vserver vs1|vs2|vs3 -destination-aggregate aggr1|aggr2"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5636168
| This role is mapped to a REST role and can only be modified by updating
the REST role.

| 5636192
| The query parameter cannot be specified for the privileges tuple with
API endpoint entries.

| 5636200
| The specified value of the access parameter is invalid, if a command or
command directory is specified in the path parameter.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#role_privilege]
[.api-collapsible-fifth-title]
role_privilege

A tuple containing a REST endpoint or a command/command directory path and
the access level assigned to that endpoint or command/command directory.
If the "path" attribute refers to a command/command directory path, the
tuple could additionally contain an optional query. The REST endpoint can
be a resource-qualified endpoint. At present, the only supported resource-
qualified endpoints are the following

Snapshots APIs

* _/api/storage/volumes/{volume.uuid}/snapshots_

File System Analytics APIs

* _/api/storage/volumes/{volume.uuid}/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/clients_
* _/api/storage/volumes/{volume.uuid}/top-metrics/directories_
* _/api/storage/volumes/{volume.uuid}/top-metrics/files_
* _/api/storage/volumes/{volume.uuid}/top-metrics/users_
* _/api/svm/svms/{svm.uuid}/top-metrics/clients_
```

```
* _/api/svm/svms/{svm.uuid}/top-metrics/directories_
* _/api/svm/svms/{svm.uuid}/top-metrics/files_
* _/api/svm/svms/{svm.uuid}/top-metrics/users_
* _/api/protocols/s3/services/{svm.uuid}/users_

In the above APIs, wildcard character * could be used in place of
_{volume.uuid}_ or _{svm.uuid}_ to denote _all_ volumes or _all_ SVMs,
depending upon whether the REST endpoint references volumes or SVMs. The
_{volume.uuid}_ refers to the _-instance-uuid_ field value in the "volume
show" command output at diagnostic privilege level. It can also be fetched
through REST endpoint _/api/storage/volumes_.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|access
|string
a|Access level for the REST endpoint or command/command directory path. If
it denotes the access level for a command/command directory path, the only
supported enum values are 'none','readonly' and 'all'.


|query
|string
a|Optional attribute that can be specified only if the "path" attribute
refers to a command/command directory path. The privilege tuple implicitly
defines a set of objects the role can or cannot access at the specified
access level. The query further reduces this set of objects to a subset of
objects that the role is allowed to access. The query attribute must be
applicable to the command/command directory specified by the "path"
attribute. It is defined using one or more parameters of the
command/command directory path specified by the "path" attribute.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SSH server

:leveloffset: +1


[[ID921bcc569ba36585297c6c3816679ec6]]
= Manage SSH server



== Overview

ONTAP supports SSH server that can be accessed from any standard SSH
client. A user account needs to be associated with SSH as the application
(refer the documentation for _api/security/accounts_
xref:{relative_path}security_accounts_endpoint_overview.html[DOC
/security/accounts] . Upon connecting from a client, the user is
authenticated and a command line shell is presented.

This endpoint is used to retrieve or modify the SSH configuration at the
cluster level. The configuration consists of SSH security parameters
(security algorithms, maximum authentication retry attempts allowed before
closing the connection, SSH connection login grace time and _ssh-rsa_
enabled status for public key algorithms) and SSH connection limits.

The security algorithms include SSH key exchange algorithms, ciphers for
payload encryption, MAC algorithms and host key algorithms. This
configuration is the default for all newly created SVMs; existing SVM
configurations are not impacted.
The SSH connection limits include maximum connections per second, maximum
simultaneous sessions from the same client host, and overall maximum SSH
connections at any given point in time. The connection limits are per node
and will be the same for all nodes in the cluster.

== Examples
```

```
=== Updating the SSH security parameters

Specify the algorithms in the body of the PATCH request.

----

# The API:
PATCH "/api/security/ssh"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/ssh" -d '{ "ciphers": [
"aes256_ctr", "aes192_ctr" ], "key_exchange_algorithms": [
"diffie_hellman_group_exchange_sha256", "ecdh_sha2_nistp256",
"diffie_hellman_group18_sha512" ], "mac_algorithms": [
"hmac_sha2_512_etm", "umac_128_etm" ], "host_key_algorithms": [
"ecdsa_sha2_nistp256", "ssh_rsa", "ssh_ed25519", "rsa_sha2_256",
"rsa_sha2_512"], "max_authentication_retry_count": 3,
"is_rsa_in_publickey_algorithms_enabled": true , "login_grace_time": 30 }'
----

=== Updating the SSH connection limits

Specify the connection limits in the body of the PATCH request.

----

# The API:
PATCH "/api/security/ssh"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/ssh" -d '{
"connections_per_second": 8, "max_instances": 10, "per_source_limit": 5 }'
----

=== Retrieving the cluster SSH server configuration

----

# The API:
GET "/api/security/ssh"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/ssh"

# The response:
{
"ciphers": [
```

```
  "aes256_ctr",
  "aes192_ctr"
],
"key_exchange_algorithms": [
  "diffie_hellman_group_exchange_sha256",
  "ecdh_sha2_nistp256",
  "diffie_hellman_group18_sha512"
],
"mac_algorithms": [
  "hmac_sha2_512_etm",
  "umac_128_etm"
],
"host_key_algorithms": [
  "ecdsa_sha2_nistp256",
  "ssh_rsa",
  "ssh_ed25519",
  "rsa_sha2_256",
  "rsa_sha2_512"
],
"max_authentication_retry_count": 3,
"connections_per_second": 8,
"max_instances": 10,
"is_rsa_in_publickey_algorithms_enabled": true,
"login_grace_time": 30,
"per_source_limit": 5,
"_links": {
  "self": {
    "href": "/api/security/ssh"
  }
}
}
----
```


```
[[ID38b48aea06b7ea298bb32df0d2e91720]]
= Retrieve cluster SSH server ciphers, MAC algorithms, key exchange
algorithms, and connection limits

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ssh`#

*Introduced In:* 9.7
```

```
Retrieves the cluster SSH server ciphers, MAC algorithms, key exchange
algorithms, host key algorithms, connection limits, login grace time, and
_ssh-rsa_ enabled status for public key algorithms.

== Related ONTAP commands

* `security ssh`
* `security protocol ssh`


== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|connections_per_second
|integer
a|Maximum connections allowed per second.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
```

```
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The SSH connection login grace time allowed for the connection.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|max_instances
|integer
a|Maximum possible simultaneous connections.


|per_source_limit
|integer
a|Maximum connections from the same client host.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "host_key_algorithms": [
    "ecdsa_sha2_nistp256",
    "ssh_rsa",
    "rsa_sha2_256",
    "rsa_sha2_512"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "ecdh_sha2_nistp256",
    "diffie_hellman_group18_sha512"
  ],
  "mac_algorithms": [
    "hmac_sha2_512",
    "hmac_sha2_512_etm"
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6a1637e2cc1d92e2049d916ea161a30d]]
= Update the SSH server setting for a cluster

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/ssh`#

*Introduced In:* 9.7

Updates the SSH server setting for a cluster.

== Optional parameters

* `ciphers` - Encryption algorithms for the payload
* `key_exchange_algorithms` - SSH key exchange algorithms
* `mac_algorithms` - MAC algorithms
* `host_key_algorithms` - Host key algorithms
* `max_authentication_retry_count` - Maximum authentication retries
allowed before closing the connection
* `connections_per_second` - Maximum allowed connections per second
* `max_instances` - Maximum allowed connections per node
* `is_rsa_in_publickey_algorithms_enabled` - _ssh-rsa_ enabled status for
public key algorithms
* `login_grace_time` - The SSH connection login grace time
* `per_source_limit` - Maximum allowed connections from the same client
```

```
host

== Related ONTAP commands

* `security ssh`
* `security protocol ssh`


== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|connections_per_second
|integer
a|Maximum connections allowed per second.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.
```

```
|login_grace_time
|integer
a|The SSH connection login grace time allowed for the connection.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|max_instances
|integer
a|Maximum possible simultaneous connections.


|per_source_limit
|integer
a|Maximum connections from the same client host.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "host_key_algorithms": [
    "ecdsa_sha2_nistp256",
    "ssh_rsa",
    "rsa_sha2_256",
    "rsa_sha2_512"
  ],
  "key_exchange_algorithms": [
```

```
    "diffie_hellman_group_exchange_sha256",
    "ecdh_sha2_nistp256",
    "diffie_hellman_group18_sha512"
  ],
  "mac_algorithms": [
    "hmac_sha2_512",
    "hmac_sha2_512_etm"
  ]
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 10682372
| There must be at least one key exchange algorithm associated with the
SSH configuration.

| 10682373
| There must be at least one cipher associated with the SSH configuration.

| 10682375
| Failed to modify SSH key exchange algorithms.

| 10682378
| Failed to modify SSH ciphers.

| 10682399
| Key exchange algorithm not supported in FIPS enabled mode.

| 10682400
| Failed to modify SSH MAC algorithms.

| 10682401
| MAC algorithm not supported in FIPS enabled mode.
```

```
| 10682403
| There must be at least one MAC algorithm with the SSH configuration.

| 10682413
| Failed to modify maximum authentication retry attempts.

| 10682413
| Failed to modify maximum authentication retry attempts.

| 10682418
| Cipher not supported in FIPS enabled mode.

| 10682423
| There must be at least one host key algorithm associated with the SSH
configuration.

| 10682424
| Host key algorithm not supported in FIPS enabled mode.

| 10682425
| Failed to modify Host key algorithms.

| 10682426
| Failed to modify _ssh-rsa_ enabled status for publickey algorithms
configuration.

| 10682426
| Failed to modify the SSH connection login grace time.

| 10682428
| Cipher not supported in FIPS enabled mode.

| 10682429
| Adding 'diffie_hellman_group16_sha512' or
'diffie_hellman_group18_sha512' to the SSH key exchange algorithms list
requires an effective cluster version of ONTAP 9.16.1 or later.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster_ssh_server]
[.api-collapsible-fifth-title]
cluster_ssh_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|connections_per_second
|integer
a|Maximum connections allowed per second.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
```

```
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The SSH connection login grace time allowed for the connection.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|max_instances
|integer
a|Maximum possible simultaneous connections.


|per_source_limit
|integer
a|Maximum connections from the same client host.


|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View SSH SVMs

:leveloffset: +1


[[ID0002e5a8d4476bbdbf51df0a2931ed64]]
= View SSH SVMs



=== Overview

This endpoint is used to retrieve the SSH security configuration for all
data SVMs. The configuration consists of SSH security parameters. The
security algorithms include SSH key exchange algorithms, ciphers for
payload encryption, MAC algorithms, host key algorithms, maximum
authentication retry attempts allowed before closing the connection, SSH
connection login grace time, and _ssh-rsa_ enabled status for public key
algorithms. The SSH configuration for a newly created SVM is the same as
the SSH configuration at cluster level. When the cluster SSH configuration
is updated using _/security/ssh_ endpoint, the SSH configuration of
existing SVMs is not impacted. To customize the SSH security parameters
for a particular data SVM, perform a PATCH operation on the
_api/security/ssh/svms/{svm.uuid}_ endpoint.

== Example

=== Retrieving the SSH security configuration of all SVMs.

Retrieves the SSH security configurations associated with the owners or
filtered list (for a specific data SVM and so on).
```

```
----

# The API:
GET "/api/security/ssh/svms"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/ssh/svms?fields=*"

# The response
{
"records": [
  {
    "svm": {
      "uuid": "739e5da5-1a50-11ee-9cb0-005056bb97dc",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/739e5da5-1a50-11ee-9cb0-005056bb97dc"
        }
      }
     },
     "ciphers": [
       "aes256_ctr",
       "aes192_ctr",
       "aes128_ctr",
       "aes128_gcm",
       "aes256_gcm"
     ],
     "key_exchange_algorithms": [
       "diffie_hellman_group_exchange_sha256",
       "ecdh_sha2_nistp256",
       "ecdh_sha2_nistp384",
       "ecdh_sha2_nistp521",
       "curve25519_sha256",
       "diffie_hellman_group16_sha512",
       "diffie_hellman_group18_sha512"
     ],
     "mac_algorithms": [
       "hmac_sha2_256",
       "hmac_sha2_512",
       "hmac_sha2_256_etm",
       "hmac_sha2_512_etm",
       "umac_64",
       "umac_128",
       "umac_64_etm",
```

```
        "umac_128_etm"
      ],
      "host_key_algorithms": [
        "ecdsa_sha2_nistp256",
        "ssh_rsa",
        "ssh_ed25519",
        "rsa_sha2_256",
        "rsa_sha2_512"
      ],
      "is_rsa_in_publickey_algorithms_enabled": true,
      "login_grace_time": 30,
      "max_authentication_retry_count": 6,
      "_links": {
        "self": {
          "href": "/api/security/ssh/svms/739e5da5-1a50-11ee-9cb0-
005056bb97dc"
        }
      }
    }
  ],
  "num_records": 1,
    "_links": {
      "self": {
        "href": "/api/security/ssh/svms?fields=*"
     }
   }
}

# The call:
curl -X GET "https://<mgmt-ip>/api/security/ssh/svms"

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    },
    "_links": {
      "self": {
```

```
        "href": "/api/security/ssh/svms/02c9e252-41be-11e9-81d5-
00a0986138f7"
      }
    }
  }
],
"num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/ssh/svms"
    }
  }
}
----
```


```
[[ID5bf2d12d309a9f40b7aefe872943a245]]
= Retrieve the SSH server configuration for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/ssh/svms`#

*Introduced In:* 9.10

Retrieves the SSH server configuration for all the data SVMs.

== Related ONTAP commands

* `security ssh`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|login_grace_time
|integer
```

```
|query
|False
a|Filter by login_grace_time

* Introduced in: 9.18


|max_authentication_retry_count
|integer
|query
|False
a|Filter by max_authentication_retry_count

* Max value: 6
* Min value: 2


|is_rsa_in_publickey_algorithms_enabled
|boolean
|query
|False
a|Filter by is_rsa_in_publickey_algorithms_enabled

* Introduced in: 9.16


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|host_key_algorithms
|string
|query
|False
a|Filter by host_key_algorithms

* Introduced in: 9.16
```

```
|mac_algorithms
|string
|query
|False
a|Filter by mac_algorithms


|key_exchange_algorithms
|string
|query
|False
a|Filter by key_exchange_algorithms


|ciphers
|string
|query
|False
a|Filter by ciphers


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#svm_ssh_server[svm_ssh_server]]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ciphers": [
        "aes256_ctr",
        "aes192_ctr",
        "aes128_ctr"
      ],
      "host_key_algorithms": [
        "ecdsa_sha2_nistp256",
        "ssh_ed25519",
        "ssh_rsa",
        "rsa_sha2_256",
        "rsa_sha2_512"
      ],
      "key_exchange_algorithms": [
        "diffie_hellman_group_exchange_sha256",
        "ecdh_sha2_nistp256"
      ],
      "mac_algorithms": [
        "hmac_sha2_512",
        "hmac_sha2_512_etm"
      ],
      "svm": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
```

```
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM name and UUID for which the SSH server is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#svm_ssh_server]
[.api-collapsible-fifth-title]
svm_ssh_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The login grace time allowed for SSH connection request timeout.
```

```
|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|svm
|link:#svm[svm]
a|SVM name and UUID for which the SSH server is configured.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SSH security configuration

:leveloffset: +1


[[ID4813f98668ffb903a1fbe5a057b95614]]
= Manage SSH security configuration



== Overview
```

```
This endpoint is used to retrieve or modify the SSH security configuration
of a data SVM.

The SSH security algorithms include key exchange algorithms, ciphers for
payload encryption, MAC algorithms, host key algorithms, SSH connection
login grace time, and the maximum authentication retry attempts allowed
before closing the connection. svm.uuid corresponds to the UUID of the SVM
for which the SSH security setting is being retrieved or modified and it
is obtained from the response body of a GET operation performed on the
_api/security/ssh/svms_ API.

== Examples

=== Updating the SSH security parameters

Specify the algorithms in the body of the PATCH request.

----

# The API:
PATCH "/api/security/ssh/svms/{svm.uuid}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/ssh/svms/02c9e252-41be-11e9-
81d5-00a0986138f7" -d '{ "ciphers": [ "aes256_ctr", "aes192_ctr" ],
"key_exchange_algorithms": [ "diffie_hellman_group_exchange_sha256",
"ecdh_sha2_nistp256", "diffie_hellman_group16_sha512" ], "mac_algorithms":
[ "hmac_sha2_512_etm", "umac_128_etm" ], "host_key_algorithms": [
"ecdsa_sha2_nistp256", "ssh_ed25519", "rsa_sha2_256", "rsa_sha2_512"
],"is_rsa_in_publickey_algorithms_enabled": false,
"max_authentication_retry_count": 3 }'
----

=== Retrieving the SSH security configuration of an SVM.

----

# The API:
GET "/api/security/ssh/svms/{svm.uuid}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/ssh/svms/02c9e252-41be-11e9-
81d5-00a0986138f7"

# The response:
{
"ciphers": [
```

```
    "aes256_ctr",
    "aes192_ctr"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "ecdh_sha2_nistp256",
    "diffie_hellman_group16_sha512"
  ],
  "mac_algorithms": [
    "hmac_sha2_512_etm",
    "umac_128_etm"
  ],
  "host_key_algorithms": [
    "ecdsa_sha2_nistp256",
    "ssh_ed25519",
    "rsa_sha2_256",
    "rsa_sha2_512"
  ],
  "is_rsa_in_publickey_algorithms_enabled": false,
  "login_grace_time": 30,
  "max_authentication_retry_count": 3,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
    "_links": {
      "self": {
        "href": "/api/svm/svms/02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/security/ssh/svms/02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
}
----
```


```
[[IDe79eeed522f4d5a3dc01b8d968db356b]]
= Retrieve the SSH server configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/security/ssh/svms/{svm.uuid}`#

*Introduced In:* 9.10

Retrieves the SSH server configuration for the specified data SVM.

== Related ONTAP commands

* `security ssh`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The login grace time allowed for SSH connection request timeout.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|svm
```

```
|link:#svm[svm]
a|SVM name and UUID for which the SSH server is configured.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "host_key_algorithms": [
    "ecdsa_sha2_nistp256",
    "ssh_ed25519",
    "ssh_rsa",
    "rsa_sha2_256",
    "rsa_sha2_512"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "ecdh_sha2_nistp256"
  ],
  "mac_algorithms": [
    "hmac_sha2_512",
    "hmac_sha2_512_etm"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
```

```
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM name and UUID for which the SSH server is configured.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDefa891668d0327c5e527a1016c868a02]]
= Update the SSH server configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/security/ssh/svms/{svm.uuid}`#

*Introduced In:* 9.10

Updates the SSH server configuration for the specified data SVM.

== Optional parameters

* `ciphers` - Encryption algorithms for the payload
* `key_exchange_algorithms` - SSH key exchange algorithms
* `host_key_algorithms` - Host key algorithms
```

```
* `mac_algorithms` - MAC algorithms
* `max_authentication_retry_count` - Maximum authentication retries
allowed before closing the connection
* `is_rsa_in_publickey_algorithms_enabled` - _ssh-rsa_ enabled status for
public key algorithms
* `login_grace_time` - Login grace time allowed for SSH connection

== Related ONTAP commands

* `security ssh`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|host_key_algorithms
```

```
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The login grace time allowed for SSH connection request timeout.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|svm
|link:#svm[svm]
a|SVM name and UUID for which the SSH server is configured.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "host_key_algorithms": [
    "ecdsa_sha2_nistp256",
    "ssh_ed25519",
    "ssh_rsa",
    "rsa_sha2_256",
    "rsa_sha2_512"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "ecdh_sha2_nistp256"
  ],
  "mac_algorithms": [
    "hmac_sha2_512",
    "hmac_sha2_512_etm"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 10682372
| There must be at least one key exchange algorithm associated with the
SSH configuration.
```

```
| 10682373
| There must be at least one cipher associated with the SSH configuration.

| 10682375
| Failed to modify SSH key exchange algorithms.

| 10682378
| Failed to modify SSH ciphers.

| 10682399
| Key exchange algorithm not supported in FIPS-enabled mode.

| 10682400
| Failed to modify SSH MAC algorithms.

| 10682401
| MAC algorithm not supported in FIPS-enabled mode.

| 10682403
| There must be at least one MAC algorithm with the SSH configuration.

| 10682413
| Failed to modify maximum authentication retry attempts.

| 10682418
| Cipher not supported in FIPS-enabled mode.

| 10682420
| To modify the SSH configuration of the admin SVM, use the
/api/security/ssh REST API.

| 10682423
| There must be at least one host key algorithm associated with the SSH
configuration.

| 10682424
| Host key algorithm not supported in FIPS enabled mode.

| 10682425
| Failed to modify Host key algorithms.

| 10682426
| Failed to modify _ssh-rsa_ enabled status for publickey algorithms
configuration.

| 10682428
| Cipher not supported in FIPS enabled mode.
```

```
| 10682429
| Adding 'diffie_hellman_group16_sha512' or
'diffie_hellman_group18_sha512' to the SSH key exchange algorithms list
requires an effective cluster version of ONTAP 9.16.1 or later.

| 10682430
| Failed to modify the login grace time.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM name and UUID for which the SSH server is configured.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#svm_ssh_server]
[.api-collapsible-fifth-title]
svm_ssh_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ciphers
|array[string]
a|Ciphers for encrypting the data.


|host_key_algorithms
|array[string]
a|Host key algorithms. The host key algorithms 'ssh_ed25519' and 'ssh_rsa'
can be configured only in non-FIPS mode.


|is_rsa_in_publickey_algorithms_enabled
|boolean
a|Enables or disables the _ssh-rsa_ signature scheme, which uses the SHA-1
hash algorithm, for RSA keys in public key algorithms. If this flag is
```

```
_false_, older SSH implementations might fail to authenticate using RSA
keys. This flag should be enabled only as a temporary measure until legacy
SSH client implementations can be upgraded or reconfigured with another
key type, for example: ECDSA.


|key_exchange_algorithms
|array[string]
a|Key exchange algorithms.


|login_grace_time
|integer
a|The login grace time allowed for SSH connection request timeout.


|mac_algorithms
|array[string]
a|MAC algorithms.


|max_authentication_retry_count
|integer
a|Maximum authentication retries allowed before closing the connection.


|svm
|link:#svm[svm]
a|SVM name and UUID for which the SSH server is configured.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage WebAuthn credentials

:leveloffset: +1


[[ID8ce3f29a98f9f32c0e8a0047383c9aac]]
= Manage WebAuthn credentials



== Overview

This API is used to retrieve all WebAuthn credentials entries.

== Examples

=== Retrieving all Webauthn credentials entries

The following output shows all WebAuthn credentials entries.

----

## The API:
/api/security/webauthn/credentials

## The call:
curl -X GET "https://<mgmt-ip>/api/security/webauthn/credentials"

# The response:
{
"records": [
  {
    "owner": {
      "uuid": "389758ee-40cd-11ef-bb21-005056aeae31",
      "_links": {
        "self": {
          "href": "/api/svm/svms/389758ee-40cd-11ef-bb21-005056aeae31"
        }
      }
    },
    "username": "user_1",
    "index": 0,
    "relying_party": {
```

```
      "id": "ontap1.example.com",
      "name": "ontap1.example.com"
    },
    "_links": {
      "self": {
        "href": "/api/security/webauthn/credentials/389758ee-40cd-11ef-
bb21-005056aeae31/user_1/0/ontap1.example.com"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/security/webauthn/credentials/"
  }
}
}
----
```


```
[[IDd4d0eba373e21cb2eadf3d2f60ed9544]]
= Retrieve all WebAuthn credentials entries

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/credentials`#

*Introduced In:* 9.16

Retrieves all WebAuthn credentials entries.

== Related ONTAP commands

* `security webauthn credentials show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#webauthn_credentials[webauthn_credentials]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "creation_time": "2024-08-05 22:38:55 -0400",
      "credential": {
        "id_sha":
"xxxxxxxxxxxxxxxxxxxxxxxebfb30f20bf6db74xxxxxxxxxxxxxxxxxxxxxxxxxx",
        "type": "string"
      },
      "display_name": "admin",
      "index": 0,
      "last_used_time": "2024-08-05 22:48:55 -0400",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "public_key": {
        "algorithm": "ES-256",
        "value":
"xxxxxxxxxxxxxxxxxxxxxe5sZohRRv2B10JjjALXmGRmEfFIxxxxxxxxxxxxxxxxxxxxxxxxx
xxxxxxxxxxxxxxxxxxxx
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#credential]
[.api-collapsible-fifth-title]
```

```
credential

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id_sha
|string
a|SHA-256 Hashed representation of Credential ID.


|type
|string
a|Credential type.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#public_key]
[.api-collapsible-fifth-title]
public_key

[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|Public key algorithm.


|value
|string
a|Public key value.


|===


[#relying_party]
[.api-collapsible-fifth-title]
```

```
relying_party

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Relying Party ID.


|name
|string
a|Relying Party Name.


|===


[#webauthn_credentials]
[.api-collapsible-fifth-title]
webauthn_credentials

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|creation_time
|string
a|Date and time indicating when this entry was created.


|credential
|link:#credential[credential]
a|

|display_name
|string
a|Display name.
```

```
|index
|integer
a|Index.


|last_used_time
|string
a|Date and time indicating when this entry was last used.


|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|public_key
|link:#public_key[public_key]
a|

|relying_party
|link:#relying_party[relying_party]
a|

|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|username
|string
a|Username.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



[[ID831690142cdca1baccf2000add33f96f]]
= Delete a WebAuthn credentials entry

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/security/webauthn/credentials/{owner.uuid}/{username}/{index}/{re
lying_party.id}`#

*Introduced In:* 9.16

Deletes a WebAuthn credentials entry.

== Required properties

***** `owner.uuid`

***** `username`

***** `index`

***** `relying_party.id`

== Related ONTAP commands

* `security webauthn credentials delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
```

```
a|Used to identify a cluster or an SVM.


|username
|string
|path
|True
a|Username.


|index
|integer
|path
|True
a|Index.


|relying_party.id
|string
|path
|True
a|Relying Party ID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
```

```
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDe88a68459d9697184630559c8f74283a]]
= Retrieve a WebAuthn credentials entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/credentials/{owner.uuid}/{username}/{index}/{re
lying_party.id}`#

*Introduced In:* 9.16

Retrieves a WebAuthn credentials entry.

== Related ONTAP commands

* `security webauthn credentials show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Used to identify a cluster or an SVM.


|username
|string
|path
|True
a|Username.


|index
|integer
|path
|True
```

```
a|Index.


|relying_party.id
|string
|path
|True
a|Relying Party ID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|creation_time
|string
a|Date and time indicating when this entry was created.


|credential
|link:#credential[credential]
a|

|display_name
|string
a|Display name.
```

```
|index
|integer
a|Index.


|last_used_time
|string
a|Date and time indicating when this entry was last used.


|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|public_key
|link:#public_key[public_key]
a|

|relying_party
|link:#relying_party[relying_party]
a|

|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|username
|string
a|Username.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```
  "creation_time": "2024-08-05 22:38:55 -0400",
  "credential": {
    "id_sha":
"xxxxxxxxxxxxxxxxxxxxxxxxebfb30f20bf6db74xxxxxxxxxxxxxxxxxxxxxxxxxx",
    "type": "string"
  },
  "display_name": "admin",
  "index": 0,
  "last_used_time": "2024-08-05 22:48:55 -0400",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "public_key": {
    "algorithm": "ES-256",
    "value":
"xxxxxxxxxxxxxxxxxxxxxe5sZohRRv2B10JjjALXmGRmEfFIxxxxxxxxxxxxxxxxxxxxxxxxx
xxxxxxxxxxxxxxxxxxxx
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#credential]
[.api-collapsible-fifth-title]
credential

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id_sha
|string
a|SHA-256 Hashed representation of Credential ID.


|type
|string
a|Credential type.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#public_key]
[.api-collapsible-fifth-title]
public_key

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|algorithm
|string
a|Public key algorithm.


|value
|string
a|Public key value.


|===


[#relying_party]
[.api-collapsible-fifth-title]
relying_party

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Relying Party ID.


|name
|string
a|Relying Party Name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View WebAuthn global settings for a cluster and all SVMs

:leveloffset: +1


[[IDc50db37aceafabd61e36de90c82d6c8c]]
= View WebAuthn global settings for a cluster and all SVMs



== Overview

This API endpoint retrieves WebAuthn global settings for a cluster and all
SVMs.

== Examples

=== Retrieving the WebAuthn global settings

Retrieves all of the WebAuthn global settings for the cluster and all
SVMs.

----

# The API:
GET "/api/security/webauthn/global-settings"

# The call to retrieve the WebAuthn global settings:
curl -k https://<mgmt-ip>/api/security/webauthn/global-settings
----





[[ID16ae550acd489903a1152d86e63b80ef]]
```

```
= Retrieve WebAuthn global settings for a cluster and all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/global-settings`#

*Introduced In:* 9.16

Retrieve WebAuthn global settings for a cluster and all SVMs.

== Related ONTAP commands

* `security webauthn show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#webauthn_global[webauthn_global]]
```

```
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "attestation": "none",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "require_rk": "",
      "resident_key": "discouraged",
      "scope": "string",
      "timeout": 600000,
      "user_verification": "discouraged"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#webauthn_global]
[.api-collapsible-fifth-title]
webauthn_global

WebAuthn global settings.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|attestation
|string
a|Attestation conveyance type.


|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|require_rk
|boolean
a|Specifies whether the resident key is required.


|resident_key
|string
a|Resident key.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.
```

```
|timeout
|integer
a|The timeout interval for the WebAuthn request, in milliseconds.


|user_verification
|string
a|User verification.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View a WebAuthn global setting entry

:leveloffset: +1


[[ID0010faa77ecb4e1687f66e07f05a285b]]
= View a WebAuthn global setting entry



== Overview

This API endpoint retrieves WebAuthn global settings for the cluster or
SVM.
Specify the owner UUID and the unique identifier for the cluster or SVM to
retrieve the WebAuthn global settings for the cluster or SVM.
```

```
== Examples

=== Retrieving the WebAuthn global settings for the cluster or SVM

----

# The API:
GET "/api/security/webauthn/global-settings/{owner.uuid}"

# The call:
curl -k https://<mgmt-ip>/api/security/webauthn/global-settings/d49de271-
8c11-11e9-8f78-005056bbf6ac
----





[[IDa07647764291a4d02061521c41ec32c3]]
= Retrieve a WebAuthn global setting entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/global-settings/{owner.uuid}`#

*Introduced In:* 9.16

Retrieves a WebAuthn global setting entry.

== Related ONTAP commands

* `security webauthn show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
```

```
|True
a|Used to identify a cluster or an SVM.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|attestation
|string
a|Attestation conveyance type.


|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|require_rk
|boolean
a|Specifies whether the resident key is required.


|resident_key
|string
a|Resident key.
```

```
|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|timeout
|integer
a|The timeout interval for the WebAuthn request, in milliseconds.


|user_verification
|string
a|User verification.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "attestation": "none",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "require_rk": "",
  "resident_key": "discouraged",
  "scope": "string",
  "timeout": 600000,
  "user_verification": "discouraged"
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1
```

```
= View all WebAuthn supported algorithms

:leveloffset: +1


[[IDc70fee19b219cde26f799090575ce974]]
= View all WebAuthn supported algorithms



== Overview

This API endpoint retrieves the list of supported algorithms for WebAuthn
authentication.

== Examples

=== Retrieving the list of supported algorithms

This endpoint retrieves all of the supported algorithms for WebAuthn
authentication. The following table lists the supported algorithms:

|===
| Algorithm ID | Algorithm Name | Description

| -7
| ES256
| ECDSA w/ SHA-256

| -35
| ES384
| ECDSA w/ SHA-384

| -257
| RS256
| RSASSA-PKCS1-v1_5 using SHA-256
|===

----

# The API:
GET "/api/security/webauthn/supported-algorithms"

# The call to retrieve the list of supported algorithms:
curl -k https://<mgmt-ip>/api/security/webauthn/supported-algorithms
```

```
----




[[ID6a50c84d67be7a05c53f289866045ee8]]
= Retrieve all WebAuthn supported algorithms entries

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/supported-algorithms`#

*Introduced In:* 9.16

Retrieve all WebAuthn supported algorithms entries.

== Related ONTAP commands

* `security webauthn supported-algorithms show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
```

```
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|
```

```
|num_records
|integer
a|Number of records.


|records
|array[link:#supported_algorithms[supported_algorithms]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "algorithm": {
        "id": -7,
        "name": "ES-256",
        "type": "string"
      },
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
```

```
      "scope": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#algorithm]
[.api-collapsible-fifth-title]
algorithm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Algorithm ID.


|name
|string
a|Algorithm name.


|type
|string
a|Algorithm type.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#supported_algorithms]
[.api-collapsible-fifth-title]
supported_algorithms

WebAuthn supported algorithms.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|algorithm
|link:#algorithm[algorithm]
a|

|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= View a WebAuthn supported algorithms entry
```

```
:leveloffset: +1


[[ID1c5fb5bdb2321db9474055ecbff5c0df]]
= View a WebAuthn supported algorithms entry



== Overview

This API endpoint retrieves a specific WebAuthn supported algorithms
entry.
Specify the owner UUID, the unique identifier for the cluster or SVM, and
the algorithm name to retrieve the specific WebAuthn supported algorithms
entry.

== Examples

=== Retrieving the specific WebAuthn supported algorithms entry

----

# The API:
GET "/api/security/webauthn/supported-algorithms/{owner.uuid}/{alg_name}"

# The call:
curl -k https://<mgmt-ip>/api/security/webauthn/supported-
algorithms/d49de271-8c11-11e9-8f78-005056bbf6ac/ES-256
----




[[ID499e0a09bb368da4c7fae27ded652e04]]
= Retrieve a WebAuthn supported algorithms entry

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/security/webauthn/supported-
algorithms/{owner.uuid}/{algorithm.name}`#

*Introduced In:* 9.16

Retrieves a WebAuthn supported algorithms entry.

== Related ONTAP commands
```

```
* `security webauthn supported-algorithms show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|Used to identify a cluster or an SVM.


|algorithm.name
|string
|path
|True
a|Algorithm name.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#self_link[self_link]
a|

|algorithm
|link:#algorithm[algorithm]
a|

|owner
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "algorithm": {
    "id": -7,
    "name": "ES-256",
    "type": "string"
  },
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
```

```
  "scope": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#algorithm]
[.api-collapsible-fifth-title]
algorithm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
```

```
a|Algorithm ID.


|name
|string
a|Algorithm name.


|type
|string
a|Algorithm type.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


:leveloffset: -1


= SnapLock

:leveloffset: +1


[[ID8b40cdcb736d234d723ba6ae45adb2c4]]
= ONTAP REST API SnapLock endpoints
```

```
SnapLock is an alternative to the traditional optical "write once, read
many" (WORM) data. SnapLock is used for the storage of read-only WORM
data. SnapLock is a license-based, disk-based, open-protocol feature that
works with application software to administer non-rewritable storage of
data. The primary objective of this feature is to provide storage-enforced
WORM and retention functionality. SnapLock can be deployed for protecting
data in strict regulatory environments in such a way that even the storage
administrator is considered an untrusted party. SnapLock provides special
purpose volumes in which files can be stored and committed to a non-
erasable, non-rewritable state either forever or for a designated
retention period. SnapLock allows this retention to be performed at the
granularity of individual files.
These APIs allow you to manage the following endpoints:

* SnapLock log
* SnapLock compliance-clock
* SnapLock event-retention
* SnapLock file retention
* SnapLock privileged-delete
* SnapLock file fingerprint
* SnapLock legal-hold

== APIs

=== SnapLock log

These APIs are used to create and initialize the SnapLock log
configuration for an SVM, modify attributes associated with the SnapLock
log configuration,and delete the auditlog configuration.

=== SnapLock compliance-clock

This API  manages the ComplianceClock of the system. ComplianceClock
determines the expiry time of the SnapLock objects in the system.
ComplianceClock can be initialized only once by the user and once it is
set, it cannot be changed or altered by the user.

=== SnapLock event-retention

This API is used to create Event Based Retention (EBR) policies for an
SVM. A policy consists of a policy-name and a retention-period. Only a
user with security login role vsadmin-snaplock is allowed to perform this
operation.

=== SnapLock file retention

This API is used to set the retention time of a file. Retention time
```

```
expects the date in ISO 8601 format or "infinite".

=== SnapLock privileged-delete

This API is used to perform a privileged-delete operation on unexpired
WORM files on a SnapLock Enterprise volume. The only built-in role that
has access to the command is "vsadmin-snaplock".

=== SnapLock file fingerprint

This API is used to start the fingerprint computation on a file.

=== SnapLock legal-hold

This API is used to retain Compliance-mode WORM files for the duration of
a litigation. A file under a Legal-Hold behaves like a WORM file with an
indefinite retention period. Only a user with security login role vsadmin-
snaplock is allowed to perform the operation.




= Manage SnapLock audit logs

:leveloffset: +1


[[ID026df54ea5741c0ff1a39a251c0e3401]]
= Manage SnapLock audit logs



The SnapLock log volume can be a SnapLock Compliance volume or SnapLock
Enterprise volume. The SnapLock log infrastructure creates directories and
files in this volume to store the SnapLock log records. The directory
named "snaplock_log" is created under the root. There are three
subdirectories created under the "snaplock_log" directory. They are
"privileged_delete_logs", "system_logs" and "legal_hold_logs". The log
files are created and populated under these subdirectories when respective
SnapLock operations are initiated. The maximum log size specifies the
maximum size of a log file that stores SnapLock log records. When the file
reaches this size, the log infrastructure archives it and creates a new
log file. An active log file has the suffix "-present" in the name. When a
file is archived, it is renamed such that the suffix is replaced with the
date of archival. The default retention period is the length of time the
log file is retained, if the SnapLock log records that are stored in the
```

```
file do not carry any retention period.

=== Examples

. Verifies that the audit log is configured for the specified SVM:

----
 GET "/api/storage/snaplock/audit-logs/?svm.name=VS0"
----

. Verifies that the specified volume is an audit log volume:

----
 GET "/api/storage/snaplock/audit-logs/?log_volume.volume.name=VS0_ALOG"
----

=== Examples

. Creates a SnapLock log configuration by providing SVM name:

----
 POST "/api/storage/snaplock/audit-logs" '{"svm": {"name":"VS3"},
"log_volume": { "volume": { "name":"VS3_ALOG"}, "max_log_size":"20971520",
"retention_period":"P30Y" }}'
----

. Creates a SnapLock log configuration by providing SVM UUID:

----
 POST "/api/storage/snaplock/audit-logs" '{"svm": {"uuid":"bc744cc7-296d-
11e9-a26f-0050568e5b05"}, "log_volume": { "volume": { "name":"VS3_ALOG"},
"max_log_size":"20971520", "retention_period":"P30Y" }}'
----

. Creates a SnapLock log configuration without specifying a retention
period:

----
 POST "/api/storage/snaplock/audit-logs" '{"svm": {"name":"VS3"},
"log_volume": {"volume": {"name":"VS3_ALOG"}}}'
----

=== Examples

. Updates the audit log volume:

----
```

```
 PATCH "/api/storage/snaplock/audit-logs/bc744cc7-296d-11e9-a26f-
0050568e5b05" '{"log_volume":{"volume":{"name":"VS4_ALOG_NEW"}}}'
----

. Updates the maximum size of the log file and the retention period:

----
 PATCH "/api/storage/snaplock/audit-logs/420cac7a-296a-11e9-a26f-
0050568e5b05" '{"log_volume":{"max_log_size":"20971520",
"retention_period":"P1Y"}}'
----

. Archives all of the audit log files:

----
 PATCH "/api/storage/snaplock/audit-logs/c7e4fa7d-2968-11e9-a26f-
0050568e5b05" '{"log_archive":{"archive":"true"}}'
----

. Archives the specified audit log file:

----
 PATCH "/api/storage/snaplock/audit-logs/c7e4fa7d-2968-11e9-a26f-
0050568e5b05"
'{"log_archive":{"archive":"true","base_name":"privileged_delete"}}'
----




[[ID6122f7be81560b0971f013bbcfe15507]]
= Retrieve SVMs configured with audit log volumes

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/audit-logs`#

*Introduced In:* 9.7

Retrieves a list of SVMs configured with audit log volumes.

== Related ONTAP commands

* `snaplock log show`

== Learn more
```

```
* xref:{relative_path}storage_snaplock_audit-
logs_endpoint_overview.html[DOC /storage/snaplock/audit-logs]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|log_files.base_name
|string
|query
|False
a|Filter by log_files.base_name


|log_files.size
|integer
|query
|False
a|Filter by log_files.size


|log_files.path
|string
|query
|False
a|Filter by log_files.path


|log_files.expiry_time
|string
|query
|False
a|Filter by log_files.expiry_time


|svm.name
|string
|query
|False
```

```
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|log_volume.max_log_size
|integer
|query
|False
a|Filter by log_volume.max_log_size


|log_volume.volume.name
|string
|query
|False
a|Filter by log_volume.volume.name


|log_volume.volume.uuid
|string
|query
|False
a|Filter by log_volume.volume.uuid


|log_volume.retention_period
|string
|query
|False
a|Filter by log_volume.retention_period


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
```

```
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snaplock_log[snaplock_log]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "log_archive": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "base_name": "string"
```

```
      },
      "log_files": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "base_name": "string",
          "expiry_time": "2058-06-04 19:00:00 +0000",
          "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
          "size": 20000
        }
      ],
      "log_volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_log_size": 20971520,
        "retention_period": "P30M",
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
```

```
a|

|===


[#log_archive]
[.api-collapsible-fifth-title]
log_archive

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|archive
|boolean
a|Archive the specified SnapLock log file for the given base_name, and
create a new log file. If base_name is not mentioned, archive all log
files.


|base_name
|string
a|Base name of log archive


|===


[#snaplock_log_file]
[.api-collapsible-fifth-title]
snaplock_log_file

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|base_name
|string
a|Base name of log file


|expiry_time
|string
a|Expiry time of the log file in date-time format. Value '9999-12-
31T00:00:00Z' indicates infinite expiry time.


|path
|string
a|Absolute path of the log file in the volume


|size
|integer
a|Size of the log file in bytes


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_log_volume]
[.api-collapsible-fifth-title]
snaplock_log_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|max_log_size
|integer
a|Maximum size of log file in bytes


|retention_period
|string
a|Specifies the default log record retention period. The retention period
value represents a duration and must be specified in the ISO-8601 duration
format. The retention period can be in years, months, days, hours, minutes
and seconds. A period specified for years, months and days is represented
in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. A duration in hours, minutes and seconds is represented by
"PT+++<num>+++H", "PT+++<num>+++M", and "PT+++<num>+++S" respectively. The
period string must contain only a single time element i.e. either years,
months, days, hours, minutes or seconds. A duration which combines
different periods is not supported, example "P1Y10M" is not supported.
Apart from the duration specified in the ISO-8601 format, the retention
period field also accepts the string
"infinite".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++
```

```
</num>+++


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snaplock_log]
[.api-collapsible-fifth-title]
snaplock_log

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|log_archive
|link:#log_archive[log_archive]
a|

|log_files
|array[link:#snaplock_log_file[snaplock_log_file]]
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID59b20597eab4f16b2fb299ac69a70505]]
= Create a SnapLock log configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/audit-logs`#
```

```
*Introduced In:* 9.7

Creates a SnapLock log configuration for an SVM. A SnapLock log
configuration consists of a volume to store the log, the maximum size of
the log file, and the default period of time for which the log file should
be retained. The volume must be a Read-Write SnapLock volume of type
Enterprise or Compliance. The input parameter retention_period expects the
duration in ISO 8601 format.

== Required properties

* `svm.uuid` or `svm.name` - Name or UUID of the SVM.
* `log_volume.volume.name` or `log_volume.volume.uuid` - Name or UUID of
audit log volume.

== Recommended optional properties

* `log_volume.max_log_size` - Max log file size.
* `log_volume.volume.retention_period` - Retention period of log file.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `log_volume.retention_period` - _P6M_
* `log_volume.max_log_size` - _10MB_

== Related ONTAP commands

* `snaplock log create`

== Learn more

* xref:{relative_path}storage_snaplock_audit-
logs_endpoint_overview.html[DOC /storage/snaplock/audit-logs]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_files
|array[link:#snaplock_log_file[snaplock_log_file]]
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "log_files": [
    {
      "base_name": "string",
      "expiry_time": "2058-06-04 19:00:00 +0000",
      "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
      "size": 20000
    }
  ],
  "log_volume": {
    "max_log_size": 20971520,
    "retention_period": "P30M",
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 918236
| The specified volume name and UUID refer to different volumes

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format

| 13763161
| Audit logging is already configured for the SVM

| 14090340
| \{field} is a required field

| 14090343
| Invalid Field

| 14090346
| Internal Error. Wait a few minutes, then try the command again
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#log_archive]
[.api-collapsible-fifth-title]
log_archive

[cols=3*,options=header]
|===
|Name
|Type
|Description

|archive
|boolean
a|Archive the specified SnapLock log file for the given base_name, and
create a new log file. If base_name is not mentioned, archive all log
files.


|base_name
|string
a|Base name of log archive


|===


[#snaplock_log_file]
[.api-collapsible-fifth-title]
snaplock_log_file

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|base_name
|string
a|Base name of log file


|expiry_time
|string
a|Expiry time of the log file in date-time format. Value '9999-12-
31T00:00:00Z' indicates infinite expiry time.


|path
|string
a|Absolute path of the log file in the volume


|size
|integer
a|Size of the log file in bytes


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_log_volume]
[.api-collapsible-fifth-title]
snaplock_log_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_log_size
|integer
a|Maximum size of log file in bytes


|retention_period
|string
a|Specifies the default log record retention period. The retention period
value represents a duration and must be specified in the ISO-8601 duration
format. The retention period can be in years, months, days, hours, minutes
and seconds. A period specified for years, months and days is represented
in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. A duration in hours, minutes and seconds is represented by
"PT+++<num>+++H", "PT+++<num>+++M", and "PT+++<num>+++S" respectively. The
period string must contain only a single time element i.e. either years,
months, days, hours, minutes or seconds. A duration which combines
different periods is not supported, example "P1Y10M" is not supported.
Apart from the duration specified in the ISO-8601 format, the retention
period field also accepts the string
"infinite".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++
</num>+++


|volume
|link:#volume[volume]
a|

|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snaplock_log]
[.api-collapsible-fifth-title]
snaplock_log

[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_files
|array[link:#snaplock_log_file[snaplock_log_file]]
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|

|svm
```

```
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd050151a6261b2892c89e65702f8ea99]]
= Disassociate SnapLock audit logs

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/audit-logs/{svm.uuid}`#

*Introduced In:* 9.7
```

```
Disassociates a SnapLock volume as the audit log volume for an SVM. This
API closes all the active log files in the log volume and marks the volume
as disabled for SnapLock logging.

== Related ONTAP commands

* `snaplock log delete`

== Learn more

* xref:{relative_path}storage_snaplock_audit-
logs_endpoint_overview.html[DOC /storage/snaplock/audit-logs]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
```

```
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090346
| Internal Error. Wait a few minutes, then try the command again
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDe13399f12c7366436cae4ed9703fb164]]
= Retrieve SnapLock audit logs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/audit-logs/{svm.uuid}`#

*Introduced In:* 9.7

Retrieves SnapLock logs for the specified SVM.

== Related ONTAP commands

* `snaplock log show`

== Learn more

* xref:{relative_path}storage_snaplock_audit-
logs_endpoint_overview.html[DOC /storage/snaplock/audit-logs]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|log_archive
|link:#log_archive[log_archive]
a|

|log_files
```

```
|array[link:#snaplock_log_file[snaplock_log_file]]
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "log_archive": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "base_name": "string"
  },
  "log_files": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "base_name": "string",
      "expiry_time": "2058-06-04 19:00:00 +0000",
      "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
      "size": 20000
    }
  ],
```

```
  "log_volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_log_size": 20971520,
    "retention_period": "P30M",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#log_archive]
[.api-collapsible-fifth-title]
log_archive

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|archive
|boolean
a|Archive the specified SnapLock log file for the given base_name, and
create a new log file. If base_name is not mentioned, archive all log
files.


|base_name
|string
a|Base name of log archive


|===


[#snaplock_log_file]
[.api-collapsible-fifth-title]
snaplock_log_file
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|base_name
|string
a|Base name of log file


|expiry_time
|string
a|Expiry time of the log file in date-time format. Value '9999-12-
31T00:00:00Z' indicates infinite expiry time.


|path
|string
a|Absolute path of the log file in the volume


|size
|integer
a|Size of the log file in bytes


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_log_volume]
[.api-collapsible-fifth-title]
snaplock_log_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|max_log_size
|integer
a|Maximum size of log file in bytes


|retention_period
|string
a|Specifies the default log record retention period. The retention period
value represents a duration and must be specified in the ISO-8601 duration
format. The retention period can be in years, months, days, hours, minutes
and seconds. A period specified for years, months and days is represented
```

```
in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. A duration in hours, minutes and seconds is represented by
"PT+++<num>+++H", "PT+++<num>+++M", and "PT+++<num>+++S" respectively. The
period string must contain only a single time element i.e. either years,
months, days, hours, minutes or seconds. A duration which combines
different periods is not supported, example "P1Y10M" is not supported.
Apart from the duration specified in the ISO-8601 format, the retention
period field also accepts the string
"infinite".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++
</num>+++


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID929d8b9581129c7090b6201029539993]]
= Update SnapLock audit logs

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/snaplock/audit-logs/{svm.uuid}`#

*Introduced In:* 9.7

Updates one of the following:

* the audit log volume,
* the attributes of the audit log volume present, or
* archive the current audit log files

== Related ONTAP commands

* `snaplock log modify`

== Learn more

* xref:{relative_path}storage_snaplock_audit-
logs_endpoint_overview.html[DOC /storage/snaplock/audit-logs]


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_archive
|link:#log_archive[log_archive]
a|

|log_files
|array[link:#snaplock_log_file[snaplock_log_file]]
```

```
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "log_archive": {
    "base_name": "string"
  },
  "log_files": [
    {
      "base_name": "string",
      "expiry_time": "2058-06-04 19:00:00 +0000",
      "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
      "size": 20000
    }
  ],
  "log_volume": {
    "max_log_size": 20971520,
    "retention_period": "P30M",
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090344
| If log_volume is specified, then log_archive must not be specified

| 14090345
| If log_archive.base_name is specified, then log_archive.archive must
also be specified
```

```
| 14090346
| Internal Error. Wait a few minutes, then try the command again
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#log_archive]
[.api-collapsible-fifth-title]
log_archive

[cols=3*,options=header]
|===
|Name
|Type
|Description

|archive
|boolean
a|Archive the specified SnapLock log file for the given base_name, and
create a new log file. If base_name is not mentioned, archive all log
files.
```

```
|base_name
|string
a|Base name of log archive


|===


[#snaplock_log_file]
[.api-collapsible-fifth-title]
snaplock_log_file

[cols=3*,options=header]
|===
|Name
|Type
|Description

|base_name
|string
a|Base name of log file


|expiry_time
|string
a|Expiry time of the log file in date-time format. Value '9999-12-
31T00:00:00Z' indicates infinite expiry time.


|path
|string
a|Absolute path of the log file in the volume


|size
|integer
a|Size of the log file in bytes


|===


[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_log_volume]
[.api-collapsible-fifth-title]
snaplock_log_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_log_size
|integer
a|Maximum size of log file in bytes


|retention_period
|string
a|Specifies the default log record retention period. The retention period
value represents a duration and must be specified in the ISO-8601 duration
format. The retention period can be in years, months, days, hours, minutes
```

```
and seconds. A period specified for years, months and days is represented
in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. A duration in hours, minutes and seconds is represented by
"PT+++<num>+++H", "PT+++<num>+++M", and "PT+++<num>+++S" respectively. The
period string must contain only a single time element i.e. either years,
months, days, hours, minutes or seconds. A duration which combines
different periods is not supported, example "P1Y10M" is not supported.
Apart from the duration specified in the ISO-8601 format, the retention
period field also accepts the string
"infinite".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++
</num>+++


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#snaplock_log]
[.api-collapsible-fifth-title]
snaplock_log

[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_archive
|link:#log_archive[log_archive]
a|

|log_files
|array[link:#snaplock_log_file[snaplock_log_file]]
a|

|log_volume
|link:#snaplock_log_volume[snaplock_log_volume]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage the SnapLock ComplianceClock

:leveloffset: +1


[[IDeb3ff33e83d6cbcee2124f9caa5784e5]]
= Manage the SnapLock ComplianceClock



This API manages the ComplianceClock of the system. ComplianceClock
determines the expiry time of the SnapLock objects in the system. The user
can initialize the ComplianceClock. The ComplianceClock can be re-
initialized if all nodes in the cluster are healthy, all volumes are in
online state, no volumes are present in the volume recovery queue and
there are no SnapLock volumes or volumes with "snapshot-locking-enabled"
set to true or S3 buckets with object locking enabled.





[[ID6c9ea409e3fcd33178d873eac25c0078]]
= Retrieve the SnapLock ComplianceClock for nodes in a cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/compliance-clocks`#

*Introduced In:* 9.7

Retrieves the SnapLock ComplianceClock for all of the nodes in the
cluster.
```

```
== Related ONTAP commands

* `snaplock compliance-clock show`

== Learn more

* xref:{relative_path}storage_snaplock_compliance-
clocks_endpoint_overview.html[DOC /storage/snaplock/compliance-clocks]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
```

```
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snaplock_compliance_clock[snaplock_compliance_clock]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "time": "2018-06-04 19:00:00 +0000"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
```

```
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#snaplock_compliance_clock]
[.api-collapsible-fifth-title]
snaplock_compliance_clock

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|node
|link:#node[node]
a|

|time
|string
a|Compliance clock time


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDfcc1d21d8a1afbbb2f589b9536fa04ef]]
= Initialize a SnapLock ComplianceClock

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/compliance-clocks`#

*Introduced In:* 9.10

Initializes the SnapLock ComplianceClock.

== Required properties

* `node.name` or `node.uuid` - Name or UUID of the node.

== Related ONTAP commands

* `snaplock compliance-clock initialize`

== Learn more

* xref:{relative_path}storage_snaplock_compliance-
clocks_endpoint_overview.html[DOC /storage/snaplock/compliance-clocks]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|time
|string
a|Compliance clock time


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "time": "2018-06-04 19:00:00 +0000"
}
```

```
====
```

```
== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763062
| Compliance Clock already initialized

| 13763084
| ComplianceClock re-initialization failed as it requires that all nodes
in the cluster are healthy, all volumes are in online state, no volumes
are present in the volume recovery queue and there are no SnapLock volumes
or volumes with "snapshot-locking-enabled" set to true or S3 buckets with
object locking enabled.

| 13763085
| ComplianceClock re-initialization supported on effective cluster version
```

```
of ONTAP 9.14.1 or later.

| 14090240
| Node with the specified UUID does not exist

| 14090241
| The specified node name and node UUID refer to different nodes

| 14090343
| Invalid field
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#snaplock_compliance_clock]
[.api-collapsible-fifth-title]
snaplock_compliance_clock

[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|time
|string
a|Compliance clock time


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDf279caa722adee8c9ddd034fc4a9c383]]
= Retrieve a SnapLock ComplianceClock for a node

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/compliance-clocks/{node.uuid}`#

*Introduced In:* 9.7

Retrieves the SnapLock ComplianceClock for a specific node.

== Related ONTAP commands

* `snaplock compliance-clock show`

== Learn more

* xref:{relative_path}storage_snaplock_compliance-
clocks_endpoint_overview.html[DOC /storage/snaplock/compliance-clocks]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|node
|link:#node[node]
a|

|time
|string
a|Compliance clock time


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "time": "2018-06-04 19:00:00 +0000"
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090240
| A node with the specified UUID does not exist

| 14090446
| Invalid session ID
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage event based retention operations

:leveloffset: +1


[[ID6557c393f6a441dece85538984d84b9e]]
= Manage event based retention operations
```

```
Use this API to display all Event Based Retention (EBR) operations and to
apply an EBR policy on a specified volume or path within the volume. Only
a user with security login role vsadmin-snaplock can perform the
operation.

=== Examples

. Displays all of the EBR operations:

----
 GET "/api/storage/snaplock/event-retention/operations"
----

. Displays all completed EBR operations:

----
 GET "/api/storage/snaplock/event-retention/operations?state=completed"
----

. Displays all completed EBR operations with filter set as volume.uuid:

----
 GET "/api/storage/snaplock/event-
retention/operations?volume.uuid=b96f976e-404b-11e9-bff2-0050568e4dbe"
----

. Displays all of the EBR operations with filter set as volume.name:

----
 GET "/api/storage/snaplock/event-retention/operations?volume.name=SLCVOL"
----

=== Examples

. Applies an EBR policy on a specific path:

----
 POST "/api/storage/snaplock/event-retention/operations"
'{"volume.name":"SLCVOL", "policy.name":"p1day", "path":"/dir1/file.txt"}'
----

. Applies an EBR policy on the complete volume:

----
 POST "/api/storage/snaplock/event-retention/operations"
'{"volume.name":"SLCVOL", "policy.name":"p1day", "path":"/"}'
```

```
----

=== Example

----
DELETE "/api/storage/snaplock/event-retention/operations/16842999"
----





[[IDfbb5272469155d1ce499d2f8d7c707e9]]
= Retrieve EBR operations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/event-retention/operations`#

*Introduced In:* 9.7

Retrieves a list of all EBR operations.

== Related ONTAP commands

* `snaplock event-retention show`

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_operations_endpoint_overview.html[DOC /storage/snaplock/event-
retention/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|query
```

```
|False
a|Volume UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
```

```
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#ebr_operation[ebr_operation]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 16842759,
      "num_files_failed": 0,
      "num_files_processed": 50,
      "num_files_skipped": 2,
      "num_inodes_ignored": 2,
      "path": "/dir1/file",
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "retention_period": "P30M"
      },
      "state": "completed",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090242
| Only a user with the security login role "vsadmin-snaplock" is allowed
to perform this operation.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#policy]
[.api-collapsible-fifth-title]
policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
```

```
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#ebr_operation]
[.api-collapsible-fifth-title]
ebr_operation

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Operation ID


|num_files_failed
|integer
```

```
a|Specifies the number of files on which the application of EBR policy
failed.


|num_files_processed
|integer
a|Specifies the number of files on which EBR policy was applied
successfully.


|num_files_skipped
|integer
a|Specifies the number of files on which the application of EBR policy was
skipped.


|num_inodes_ignored
|integer
a|Specifies the number of inodes on which the application of EBR policy
was not attempted because they were not regular files.


|path
|string
a|The path for the EBR operation. Specifies the path relative to the
output volume root, of the form "/path". The path can be path to a file or
a directory.


|policy
|link:#policy[policy]
a|

|state
|string
a|Specifies the operation status of an EBR operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID52448216927d1fc2f043c63b9ae05e90]]
= Create an EBR policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/event-retention/operations`#

*Introduced In:* 9.7

Start an EBR operation. This operation applies an EBR policy to a specific
volume or to a path within the volume.

== Required properties

* `path` - Path of the file.
* `policy.name` - Name of the EBR policy.
Either volume.name or volume.uuid must be specified when applying EBR
policy to the volume or a path within the volume.
* `volume.name` - Volume name.
* `volume.uuid` - Volume uuid.

== Related ONTAP commands

* `snaplock event-retention apply`

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_operations_endpoint_overview.html[DOC /storage/snaplock/event-
retention/operations]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID


|num_files_failed
|integer
a|Specifies the number of files on which the application of EBR policy
failed.


|num_files_processed
|integer
a|Specifies the number of files on which EBR policy was applied
successfully.
```

```
|num_files_skipped
|integer
a|Specifies the number of files on which the application of EBR policy was
skipped.


|num_inodes_ignored
|integer
a|Specifies the number of inodes on which the application of EBR policy
was not attempted because they were not regular files.


|path
|string
a|The path for the EBR operation. Specifies the path relative to the
output volume root, of the form "/path". The path can be path to a file or
a directory.


|policy
|link:#policy[policy]
a|

|state
|string
a|Specifies the operation status of an EBR operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 50,
  "num_files_skipped": 2,
  "num_inodes_ignored": 2,
  "path": "/dir1/file",
  "policy": {
    "name": "string",
    "retention_period": "P30M"
  },
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID


|num_files_failed
|integer
a|Specifies the number of files on which the application of EBR policy
failed.


|num_files_processed
|integer
a|Specifies the number of files on which EBR policy was applied
```

```
successfully.


|num_files_skipped
|integer
a|Specifies the number of files on which the application of EBR policy was
skipped.


|num_inodes_ignored
|integer
a|Specifies the number of inodes on which the application of EBR policy
was not attempted because they were not regular files.


|path
|string
a|The path for the EBR operation. Specifies the path relative to the
output volume root, of the form "/path". The path can be path to a file or
a directory.


|policy
|link:#policy[policy]
a|

|state
|string
a|Specifies the operation status of an EBR operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 50,
  "num_files_skipped": 2,
  "num_inodes_ignored": 2,
  "path": "/dir1/file",
  "policy": {
    "name": "string",
    "retention_period": "P30M"
  },
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```


Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 918236
| The specified volume uuid and volume name refer to different volumes.

| 13763372
| Path should be given in the format "\/<dir path>".

| 14090242
| Only a user with the security login role "vsadmin-snaplock" is allowed
to perform this operation.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#policy]
[.api-collapsible-fifth-title]
policy
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#ebr_operation]
[.api-collapsible-fifth-title]
ebr_operation
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID


|num_files_failed
|integer
a|Specifies the number of files on which the application of EBR policy
failed.


|num_files_processed
|integer
a|Specifies the number of files on which EBR policy was applied
successfully.


|num_files_skipped
|integer
a|Specifies the number of files on which the application of EBR policy was
skipped.


|num_inodes_ignored
|integer
a|Specifies the number of inodes on which the application of EBR policy
was not attempted because they were not regular files.


|path
|string
a|The path for the EBR operation. Specifies the path relative to the
output volume root, of the form "/path". The path can be path to a file or
a directory.


|policy
|link:#policy[policy]
a|

|state
```

```
|string
a|Specifies the operation status of an EBR operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDe6e30ebb0c3baa65f7da2377a7c65521]]
= Cancel an ongoing EBR operation

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/event-retention/operations/{id}`#

*Introduced In:* 9.7

Aborts an ongoing EBR operation.

== Related ONTAP commands

* `snaplock event-retention abort`

== Learn more

* xref:{relative_path}storage_snaplock_event-
```

```
retention_operations_endpoint_overview.html[DOC /storage/snaplock/event-
retention/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Unique identifier of an EBR operation

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090541
| A completed or failed operation cannot be aborted

| 14090242
| Only a user with the security login role "vsadmin-snaplock" is allowed
to perform this operation.
|===
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID862512ad303e8c4090f86fd78b49af73]]
= Retrieve attributes for an EBR operation

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/event-retention/operations/{id}`#

*Introduced In:* 9.7

Retrieves a list of attributes for an EBR operation.

== Related ONTAP commands

* `snaplock event-retention show`

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_operations_endpoint_overview.html[DOC /storage/snaplock/event-
retention/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Unique identifier of an EBR operation
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Operation ID


|num_files_failed
|integer
a|Specifies the number of files on which the application of EBR policy
failed.


|num_files_processed
|integer
a|Specifies the number of files on which EBR policy was applied
successfully.


|num_files_skipped
|integer
a|Specifies the number of files on which the application of EBR policy was
skipped.


|num_inodes_ignored
```

```
|integer
a|Specifies the number of inodes on which the application of EBR policy
was not attempted because they were not regular files.


|path
|string
a|The path for the EBR operation. Specifies the path relative to the
output volume root, of the form "/path". The path can be path to a file or
a directory.


|policy
|link:#policy[policy]
a|

|state
|string
a|Specifies the operation status of an EBR operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 50,
  "num_files_skipped": 2,
```

```
  "num_inodes_ignored": 2,
  "path": "/dir1/file",
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string",
    "retention_period": "P30M"
  },
  "state": "completed",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```


Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090242
| Only a user with the security login role "vsadmin-snaplock" is allowed
to perform this operation.
```

```
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#policy]
[.api-collapsible-fifth-title]
policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Specifies the EBR policy name
```

```
|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage event based retention policies

:leveloffset: +1


[[ID02a3887a7745a7441a1e88b02d60675f]]
= Manage event based retention policies



An event retention policy consists of a policy-name and a retention-
period. The policy can be applied to a single file or files in a
directory. Only a user with the security login role vsadmin-snaplock can
perform the operation. EBR policies cannot be applied to files under a
Legal-Hold.

=== Examples

. Creates an EBR policy policy_name with a retention period of "10 years":

----
 POST "/api/storage/snaplock/event-retention/policies/" '{"name":
"policy_name","retention_period": "P10Y"}'
----

. Creates an EBR policy policy_name1 with a retention period of
"infinite":

----
 POST "/api/storage/snaplock/event-retention/policies/" '{"name":
"policy_name1","retention_period": "infinite"}'
----
```

```
[[IDa78639b5c7037951123185ca2daaf4b8]]
= Retrieve event retention policies for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/event-retention/policies`#

*Introduced In:* 9.7

Retrieves all event retention policies for an SVM.

== Related ONTAP commands

* `snaplock event-retention policy show`

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_policies_endpoint_overview.html[DOC /storage/snaplock/event-
retention/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#snaplock_retention_policy[snaplock_retention_policy]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "retention_period": "P30M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
```

```
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format or infinite or unspecified.

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
```

```
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
```

```
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snaplock_retention_policy]
[.api-collapsible-fifth-title]
snaplock_retention_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID3d0c58e4701c09caba21a0a5b9fdb1a8]]
= Create an EBR policy for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/event-retention/policies`#

*Introduced In:* 9.7

Creates an Event Based Retention (EBR) policy for an SVM. The input
parameter retention_period expects the duration in ISO 8601 format or
infinite.

== Required properties

* `name` - Event retention policy name.
* `retention_period` - Retention period of the EBR policy.

== Related ONTAP commands

* `snaplock event-retention policy create`

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_policies_endpoint_overview.html[DOC /storage/snaplock/event-
```

```
retention/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
```

```
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the EBR policy name
```

```
|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
```

```
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format or infinite or unspecified.

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 13763402
| Policy name already exists.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snaplock_retention_policy]
[.api-collapsible-fifth-title]
snaplock_retention_policy
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[ID6eabb1fd552ff21bef43f58877cde9a2]]
= Delete an EBR policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/event-retention/policies/{policy.name}`#

*Introduced In:* 9.7

Deletes the specified Event Based Retention (EBR) policy.

== Related ONTAP commands

* `snaplock event-retention policy delete`

== Example

Delete the specified Event Based Retention policy "policy_name":

----
DELETE "/api/storage/snaplock/event-retention/policies/{policy.name}"
----

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_policies_endpoint_overview.html[DOC /storage/snaplock/event-
retention/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|policy.name
|string
|path
```

```
|True
a|Name of the retention policy

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID38b01e0a8c81e1dc43ed991d5c26d77e]]
= Retrieve attributes of an EBR policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/event-retention/policies/{policy.name}`#

*Introduced In:* 9.7

Retrieves a list of attributes of the specified Event Based Retention
(EBR) policy.

== Related ONTAP commands

* `snaplock event-retention policy show`
```

```
== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_policies_endpoint_overview.html[DOC /storage/snaplock/event-
retention/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|policy.name
|string
|path
|True
a|Name of the retention policy


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|name
|string
a|Specifies the EBR policy name


|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "retention_period": "P30M",
  "svm": {
```

```
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



[[ID4954c81eba0fd0b096d69cb6442724aa]]
= Update the retention period of an EBR policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/snaplock/event-retention/policies/{policy.name}`#

*Introduced In:* 9.7

Updates the retention period of an Event Based Retention (EBR) policy.

== Related ONTAP commands

* `snaplock event-retention policy modify`

== Example

Updates the retention period of an EBR policy "policy_name":

----
PATCH "/api/storage/snaplock/event-retention/policies/{policy.name}"
'{"retention_period": "P20Y"}'
----

== Learn more

* xref:{relative_path}storage_snaplock_event-
retention_policies_endpoint_overview.html[DOC /storage/snaplock/event-
retention/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|policy.name
```

```
|string
|path
|True
a|Name of the retention policy

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format or infinite or unspecified.

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snaplock_retention_policy]
[.api-collapsible-fifth-title]
snaplock_retention_policy

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|retention_period
|string
a|Specifies the retention period of an event based retention policy. The
retention period value represents a duration and must be specified in the
ISO-8601 duration format. The retention period can be in years, months,
days, hours or minutes. A period specified for years, months and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively. For example "P10Y" represents a duration of
10 years. Similarly, a duration in hours, minutes is represented by
"PT+++<num>+++H", "PT+++<num>+++M" respectively. The period string must
contain only a single time element i.e. either years, months, days, hours
or minutes. A duration which combines different periods is not supported,
example "P1Y10M" is not supported. Apart from the duration specified in
the ISO-8601 format, the retention period field also accepts the strings
"infinite" and
"unspecified".+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1
```

```
= View key information about files and volumes

:leveloffset: +1


[[ID15b6ed939b99c74b3a60b04f5c5aba50]]
= View key information about files and volumes



Use this API to view key information about files and volumes, including
the file type (regular, WORM, or WORM appendable), the volume expiration
date, and so on.





[[IDb92ee3dc152ab3df99dcb5c733a82ec9]]
= Retrieve fingerprint operations of an SVM and volume

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/file-fingerprints`#

*Introduced In:* 9.7

Retrieves a list of all the fingerprint operations of the specified SVM
and volume.

== Related ONTAP commands

* `volume file fingerprint show`

== Example

----
GET "/api/storage/snaplock/file-fingerprints/?svm.uuid=23940494-3f3a-11e9-
8675-0050568e8f89&volume.uuid=36cdb58c-3f3a-11e9-8675-0050568e8f89"
----

== Learn more

* xref:{relative_path}storage_snaplock_file-
fingerprints_endpoint_overview.html[DOC /storage/snaplock/file-
fingerprints]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|query
|True
a|SVM UUID


|volume.uuid
|string
|query
|True
a|Volume UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of Records.


|records
|array[link:#snaplock_file_fingerprint_get[snaplock_file_fingerprint_get]]
```

```
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 17039367,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
```

```
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090442
| Invalid Vserver UUID

| 14090445
| Invalid key values. Provide Vserver UUID and Volume UUID
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
```

```
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
```

```
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_file_fingerprint_get]
[.api-collapsible-fifth-title]
snaplock_file_fingerprint_get

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|A unique identifier for the fingerprint operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDcb3d39180dd0926905341c9439e4247e]]
= Create a fingerprint computation session

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/file-fingerprints`#

*Introduced In:* 9.7
```

```
Creates a fingerprint computation session on the file and returns a
session-id. This session-id is a unique identifier that you can use to
retrieve the progress of an ongoing fingerprint operation. When the
operation is complete, you can use the session-id to retrieve the complete
fingerprint output for the file .

== Required properties

* `svm.uuid` or `svm.name` - Name or UUID of the SVM.
* `volume.name` or `volume.uuid` - Name or UUID of the volume.
* `path` - Path of the file.

== Default property values

If not specified in POST, the follow default property values are assigned:

* `algorithm` - _sha256_

== Related ONTAP commands

* `volume file fingerprint start`

== Example

----
POST "/api/storage/snaplock/file-fingerprints" '{"svm":{"uuid":"23940494-
3f3a-11e9-8675-0050568e8f89"},"volume": {"uuid":"26cdb58c-3f3a-11e9-8675-
0050568e8f89"},"path":"/vol/a1.txt","algorithm":"md5"}'
----

== Learn more

* xref:{relative_path}storage_snaplock_file-
fingerprints_endpoint_overview.html[DOC /storage/snaplock/file-
fingerprints]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|The digest algorithm which is used for the fingerprint computation


|data_fingerprint
|string
a|The digest value of data of the file. The fingerprint is base64 encoded.
This field is not included if the scope is metadata-only.


|file_size
|integer
a|The size of the file in bytes.


|file_type
|string
a|The type of the file.


|id
|integer
a|A unique identifier for the fingerprint operation
```

```
|metadata_fingerprint
|string
a|The digest value of metadata of the file. The metadata fingerprint is
calculated for file size, file ctime, file mtime, file crtime, file
retention time, file uid, file gid, and file type. The fingerprint is
base64 encoded. This field is not included if the scope is data-only.


|path
|string
a|Specifies the path on which file fingerprint operation is running or has
completed. Specifies the path relative to the output volume root, of the
form "/path". The path can be path to a file or a directory.


|scope
|string
a|The scope of the file which is used for the fingerprint computation


|state
|string
a|Specifies the status of fingerprint operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "algorithm": "sha256",
  "data_fingerprint": "MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=",
  "file_size": 1048576,
  "file_type": "worm",
  "id": 17039367,
```

```
  "metadata_fingerprint": "8iMjqJXiNcqgXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=",
  "path": "/homedir/dir1",
  "scope": "data_only",
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```


Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description
```

```
| 2
| An invalid value was entered for one of the fields

| 2621519
| Invalid SVM name format

| 14090347
| Invalid path format

| 14090443
| Invalid volume name

| 14090444
| Invalid Vserver name

| 14090447
| Invalid volume UUID

| 14090448
| Invalid key values. Provide valid Vserver name and volume name or
Vserver UUID and volume UUID
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_file_fingerprint]
[.api-collapsible-fifth-title]
snaplock_file_fingerprint

[cols=3*,options=header]
|===
|Name
|Type
|Description

|algorithm
|string
a|The digest algorithm which is used for the fingerprint computation


|data_fingerprint
|string
a|The digest value of data of the file. The fingerprint is base64 encoded.
This field is not included if the scope is metadata-only.


|file_size
|integer
a|The size of the file in bytes.


|file_type
|string
a|The type of the file.


|id
```

```
|integer
a|A unique identifier for the fingerprint operation


|metadata_fingerprint
|string
a|The digest value of metadata of the file. The metadata fingerprint is
calculated for file size, file ctime, file mtime, file crtime, file
retention time, file uid, file gid, and file type. The fingerprint is
base64 encoded. This field is not included if the scope is data-only.


|path
|string
a|Specifies the path on which file fingerprint operation is running or has
completed. Specifies the path relative to the output volume root, of the
form "/path". The path can be path to a file or a directory.


|scope
|string
a|The scope of the file which is used for the fingerprint computation


|state
|string
a|Specifies the status of fingerprint operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID05186080e352e9c0b9ecb185f6c1293f]]
= Cancel a fingerprint operation

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/file-fingerprints/{id}`#

*Introduced In:* 9.7

Aborts an in-progress fingerprint operation. This API takes session-id as
input and aborts the fingerprint operation that is associated with the
specified session-id.

== Related ONTAP commands

* `volume file fingerprint abort`

== Learn more

* xref:{relative_path}storage_snaplock_file-
fingerprints_endpoint_overview.html[DOC /storage/snaplock/file-
fingerprints]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Unique identifier of the operation
```

```
|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090440
| File fingerprint operation has completed

| 14090446
| Invalid session ID
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID7d39d83e173fd5e15c295369eafd42c0]]
= Retrieve the file fingerprint information for a specific session ID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/file-fingerprints/{id}`#

*Introduced In:* 9.7

Retrieves the file fingerprint information for a specific session ID.

== Related ONTAP commands

* `volume file fingerprint dump`
```

```
== Learn more

* xref:{relative_path}storage_snaplock_file-
fingerprints_endpoint_overview.html[DOC /storage/snaplock/file-
fingerprints]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Unique identifier of the operation


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|algorithm
|string
a|The digest algorithm which is used for the fingerprint computation


|data_fingerprint
|string
a|The digest value of data of the file. The fingerprint is base64 encoded.
This field is not included if the scope is metadata-only.


|file_size
|integer
a|The size of the file in bytes.


|file_type
|string
a|The type of the file.


|id
|integer
a|A unique identifier for the fingerprint operation


|metadata_fingerprint
|string
a|The digest value of metadata of the file. The metadata fingerprint is
calculated for file size, file ctime, file mtime, file crtime, file
retention time, file uid, file gid, and file type. The fingerprint is
base64 encoded. This field is not included if the scope is data-only.


|path
|string
a|Specifies the path on which file fingerprint operation is running or has
completed. Specifies the path relative to the output volume root, of the
form "/path". The path can be path to a file or a directory.


|scope
|string
a|The scope of the file which is used for the fingerprint computation
```

```
|state
|string
a|Specifies the status of fingerprint operation.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "algorithm": "sha256",
  "data_fingerprint": "MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=",
  "file_size": 1048576,
  "file_type": "worm",
  "id": 17039367,
  "metadata_fingerprint": "8iMjqJXiNcqgXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=",
  "path": "/homedir/dir1",
  "scope": "data_only",
  "state": "completed",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
```

```
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 14090441
| Invalid key values. Provide session ID or Vserver UUID and volume UUID

| 14090449
| File fingerprint operation has not completed
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SnapLock file retention time
```

```
:leveloffset: +1


[[IDa68ae4dee55432f81843e0f70573f29f]]
= Manage SnapLock file retention time



This API manages the SnapLock retention time of a file. You can perform a
privileged-delete operation by executing this API.

=== Examples

. Sets the SnapLock retention time of a file:

----
 PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-
0050568e3f82/%2Ffile2.txt" '{"expiry_time": "2030-02-14T18:30:00+5:30"}'
----

. Extends the retention time of a WORM file:

----
 PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-
0050568e3f82/%2Ffile2.txt" '{"expiry_time": "infinite"}'
----

. Extends the retention time of a WORM file:

----
 PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-
0050568e3f82/%2Ffile2.txt" '{"retention_period": "P1M"}'
----

. Extends the retention time of a WORM file:

----
 PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-
0050568e3f82/%2Ffile2.txt" '{"retention_period": "infinite"}'
----





[[ID4dae88cce678645065e58a508e5865d6]]
= Delete unexpired WORM files of a SnapLock enterprise volume
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/file/{volume.uuid}/{path}`#

*Introduced In:* 9.7

Deletes unexpired WORM files of a SnapLock Enterprise volume. This is a
privileged-delete operation. The only built-in role that has access to the
command is vsadmin-snaplock.

== Related ONTAP commands

* `volume file privileged-delete`

== Learn more

*
xref:{relative_path}storage_snaplock_file_volume.uuid_path_endpoint_overvi
ew.html[DOC /storage/snaplock/file/{volume.uuid}/\{path}]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|path
|string
|path
|True
a|Path of the file in the form "/<dirpath>/<filename>"


|return_timeout
```

```
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 262179
| Unexpected argument "+++<file_name>+++\"+++</file_name>+++
```

```
| 6691623
| User is not authorized

| 13763162
| SnapLock audit log volume is not configured for the SVM

| 13763280
| Only a user with the security login role "vsadmin-snaplock" is allowed
to perform this operation

| 14090347
| File path must be in the format "\/<dir>\/<file path>"

| 917804
| Path should be given in the format "\/\vol\/<volume name>\/<file path>".
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
```

```
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
```

```
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID66e99338f5f6bf10c45e1300f12ebd2a]]
= Retrieve SnapLock retention details of a file

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/file/{volume.uuid}/{path}`#

*Introduced In:* 9.7

Retrieves the SnapLock retention details of the specified file. An
indefinite expiry time indicates the file is under a Legal-Hold.
```

```
== Related ONTAP commands

* `volume file retention show`

== Learn more

*
xref:{relative_path}storage_snaplock_file_volume.uuid_path_endpoint_overvi
ew.html[DOC /storage/snaplock/file/{volume.uuid}/\{path}]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|path
|string
|path
|True
a|Path of the file in the form "/<dirpath>/<filename>"


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|expiry_time
|string
a|Expiry time of the file in date-time format, "infinite", "indefinite",
or "unspecified". An "infinite" retention time indicates that the file
will be retained forever. An "unspecified" retention time indicates that
the file will be retained forever; however, the retention time of the file
can be changed to an absolute value. An "indefinite" retention time
indicates that the file is under Legal-Hold.


|file_path
|string
a|Specifies the volume relative path of the file


|is_expired
|boolean
a|This indicates if the file is under active retention or if the file is
past its expiry time.


|seconds_until_expiry
|integer
a|Specifies the number of seconds until the expiration time of the file.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "expiry_time": "2058-06-04 19:00:00 +0000",
  "file_path": "/dir1/file",
  "is_expired": 1,
  "seconds_until_expiry": 168,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes
```

```
|===
| Error code | Description

| 262179
| Unexpected argument "+++<file_name>+++\"+++</file_name>+++

| 917864
| Operation not supported

| 6691623
| User is not authorized

| 13762592
| Operation not supported on non-SnapLock volume

| 14090347
| File path must be in the format "\/<dir>\/<file path>"

| 917804
| Path should be given in the format "\/\vol\/<volume name>\/<file path>".
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
```

```
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0d77c431c20da7bc491ea020c6918943]]
= Update the SnapLock retention time of a file

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/snaplock/file/{volume.uuid}/{path}`#

*Introduced In:* 9.7

Updates the SnapLock retention time of a file or extends the retention
time of a WORM file. Input parameters "expiry_time" or "retention_period"
can be used to set or extend the retention time of file. Both
"expiry_time" and "retention_period" parameters expect the date in ISO
```

```
8601 format. Additionally, the "expiry_time" parameter can also be set to
"infinite" or "unspecified" and the "retention_period" parameter can also
be set to "infinite". The input parameters are mutually exclusive.

== Related ONTAP commands

* `volume file retention set`

== Learn more

*
xref:{relative_path}storage_snaplock_file_volume.uuid_path_endpoint_overvi
ew.html[DOC /storage/snaplock/file/{volume.uuid}/\{path}]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|path
|string
|path
|True
a|Path of the file in the form "/<dirpath>/<filename>"

|===

== Request Body



[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|expiry_time
|string
a|Expiry time of the file in date-time format, "infinite", "indefinite",
or "unspecified". An "infinite" retention time indicates that the file
will be retained forever. An "unspecified" retention time indicates that
the file will be retained forever; however, the retention time of the file
can be changed to an absolute value. An "indefinite" retention time
indicates that the file is under Legal-Hold.


|is_expired
|boolean
a|This indicates if the file is under active retention or if the file is
past its expiry time.


|retention_period
|string
a|Duration of retention time file to be locked with,  An "infinite"
retention period indicates that the file will be retained forever.


|seconds_until_expiry
|integer
a|Specifies the number of seconds until the expiration time of the file.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "expiry_time": "2058-06-04 19:00:00 +0000",
  "is_expired": 1,
  "retention_period": "P2M",
  "seconds_until_expiry": 168
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 262179
| Unexpected argument "+++<file_name>+++\"+++</file_name>+++

| 262186
| Field "expiry_time" cannot be used with field "retention_period"

| 6691623
| User is not authorized

| 13763279
| The resulting expiry time due to the specified retention period is
earlier than the current expiry time

| 14090348
| Invalid Expiry time

| 14090347
| File path must be in the format "\/<dir>\/<file path>"

| 917804
| Path should be given in the format "\/\vol\/<volume name>\/<file path>".
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_file_retention]
[.api-collapsible-fifth-title]
snaplock_file_retention

[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|string
a|Expiry time of the file in date-time format, "infinite", "indefinite",
or "unspecified". An "infinite" retention time indicates that the file
will be retained forever. An "unspecified" retention time indicates that
the file will be retained forever; however, the retention time of the file
can be changed to an absolute value. An "indefinite" retention time
indicates that the file is under Legal-Hold.
```

```
|is_expired
|boolean
a|This indicates if the file is under active retention or if the file is
past its expiry time.


|retention_period
|string
a|Duration of retention time file to be locked with,  An "infinite"
retention period indicates that the file will be retained forever.


|seconds_until_expiry
|integer
a|Specifies the number of seconds until the expiration time of the file.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SnapLock storage litigations

:leveloffset: +1


[[IDbf95834330bcc051e8a5f61bd11a6f6d]]
= Manage SnapLock storage litigations
```

```
Use this API to retain Compliance-mode WORM files for the duration of a
litigation. A file under a legal-hold behaves as a WORM file with an
indefinite retention period. Litigation ID is a combination of volume UUID
and litigation name in the format `<volume UUID>:<litigation name>`. Only
a user with the security login role vsadmin-snaplock can perform the
operation.





[[IDb25664876d936376172f48e693f6b4b1]]
= Retrieve litigations under an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/litigations`#

*Introduced In:* 9.7

Retrieves the list of litigations under an SVM.

== Related ONTAP commands

* `snaplock legal-hold show`

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_endpoint_overview.html[DO
C /storage/snaplock/litigations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
|query
|False
```

```
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of Records


|records
|array[link:#snaplock_litigation_get[snaplock_litigation_get]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "string",
      "name": "lit1",
      "operations": [
```

```
        {
          "id": 16842759
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090346
| Internal Error. Wait a few minutes, then try the command again

| 14090343
| Invalid field
```

```
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#operations]
[.api-collapsible-fifth-title]
operations

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
```

```
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#snaplock_litigation_get]
[.api-collapsible-fifth-title]
snaplock_litigation_get

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|string
a|Specifies the litigation ID.


|name
|string
a|Specifies the legal-hold litigation name.


|operations
|array[link:#operations[operations]]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDa1bbe9c165422d81e1b781e34fc74d7e]]
= Start a legal-hold operation

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/litigations`#

*Introduced In:* 9.7

Starts a  Legal-Hold.

== Required properties

* `path` - Path of the file.
* `name` - Litigation name.
* `volume.name` or `volume.uuid` - Name or UUID  of the volume.

== Related ONTAP commands

* `snaplock legal-hold begin`

== Example

----
POST "/api/storage/snaplock/litigations"
'{"volume.name":"SLC1","name":"l3","path":"/b.txt"}'
----

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_endpoint_overview.html[DO
C /storage/snaplock/litigations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
```

```
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Specifies the litigation ID.


|name
|string
a|Specifies the legal-hold litigation name.


|operations
|array[link:#snaplock_legal_hold_operation[snaplock_legal_hold_operation]]
a|

|path
|string
a|Specifies the path on which legal-hold operation has to be applied.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "id": "string",
  "name": "lit1",
  "operations": [
    {
      "id": 16842759,
      "num_files_failed": 0,
      "num_files_processed": 30,
      "num_files_skipped": 10,
      "num_inodes_ignored": 10,
      "path": "/dir1",
      "state": "completed",
      "type": "begin"
    }
  ],
  "path": "/dir1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13762586
| Volume is not a SnapLock Compliance volume

| 13763162
| Audit logging is not configured for the SVM

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090346
| Internal Error. Wait a few minutes, then try the command again

| 14090340
| \{field} is a required field

| 14090343
| Invalid Field
```

```
| 14090641
| The specified volume.name and volume.uuid refer to different volumes
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#snaplock_legal_hold_operation]
[.api-collapsible-fifth-title]
snaplock_legal_hold_operation

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID.
```

```
|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.


|num_files_processed
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.


|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.


|type
|string
a|Specifies the type of legal-hold operation.


|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
* x-nullable: true


|===


[#snaplock_litigation]
[.api-collapsible-fifth-title]
snaplock_litigation

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Specifies the litigation ID.


|name
|string
a|Specifies the legal-hold litigation name.


|operations
|array[link:#snaplock_legal_hold_operation[snaplock_legal_hold_operation]]
a|

|path
|string
a|Specifies the path on which legal-hold operation has to be applied.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID1958224c9b1a9a66228b5121c4b8d160]]
= Create a legal-hold end on all of the files for a litigation ID

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/litigations/{id}`#

*Introduced In:* 9.7

Creates a legal-hold end on all of the files for the specified litigation
ID. This is only allowed when an operation is no longer in progress.

== Related ONTAP commands

* `snaplock legal-hold end`

== Example

----
DELETE "/api/storage/snaplock/litigations/fd72e138-4bc3-11e9-a85f-
0050568eb48f%3Al3"
----

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_endpoint_overview.html[DO
C /storage/snaplock/litigations]


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Litigation ID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090343
| Invalid Field

| 14090346
| Internal Error. Wait a few minutes, then try the command again

| 14090642
| Legal-Hold operation is in progress. Legal-Hold end cannot be performed.
Wait for the operation to complete or abort the operation
|===
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



[[ID2fc871c15264d8f68f888fc8e2115e24]]
= Retrieve ongoing operations for a litigation ID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/litigations/{id}`#

*Introduced In:* 9.7

Retrieves the list of ongoing operations for the specified litigation ID.

== Related ONTAP commands

* `snaplock legal-hold show`

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_endpoint_overview.html[DO
C /storage/snaplock/litigations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Litigation ID


|fields
```

```
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|string
a|Specifies the litigation ID.


|name
|string
a|Specifies the legal-hold litigation name.


|operations
|array[link:#snaplock_legal_hold_operation[snaplock_legal_hold_operation]]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": "string",
  "name": "lit1",
  "operations": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 16842759,
      "num_files_failed": 0,
      "num_files_processed": 30,
      "num_files_skipped": 10,
      "num_inodes_ignored": 10,
      "path": "/dir1",
      "state": "completed",
      "type": "begin"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
```

```
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090346
| Internal Error. Wait a few minutes, then try the command again

| 14090343
| Invalid Field
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
```

```
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#snaplock_legal_hold_operation]
[.api-collapsible-fifth-title]
snaplock_legal_hold_operation

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Operation ID.


|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.


|num_files_processed
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.


|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
```

```
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.


|type
|string
a|Specifies the type of legal-hold operation.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= View litigation ID files

:leveloffset: +1


[[ID1ea7af9e17c1e65ed2d24e8b74da70fc]]
= View litigation ID files



Displays the list of files under the specified litigation ID.





[[IDa4f3922f0382241c1b21341240df6827]]
= Display files for a litigation ID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/litigations/{litigation.id}/files`#

*Introduced In:* 9.7

Displays the list of files for the specified litigation ID.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|litigation.id
|string
|path
|True
a|Litigation ID


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
```

```
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|List of Files under the specified litigation.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
```

```
  },
  "num_records": 1,
  "records": [
    {
      "file": [
        "string"
      ]
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|file
|array[string]
a|Name of the file including the path from the root.


|sequence_index
|integer
a|Sequence index of files path list.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage legal-hold operations

:leveloffset: +1


[[ID5efeb1f4617a24487dbc792ddfd97555]]
= Manage legal-hold operations



Manages the legal-hold operations for the specified litigation ID.

=== Examples

. Adds a Legal-Hold.

----
 POST "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-
0050568ebef5:l1/operations" '{"type" : "begin", "path" : "/a.txt"}'
----

. Removes a Legal-Hold.

----
 POST "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-
0050568ebef5:l1/operations" '{"type" : "end", "path" : "/a.txt"}'
----





[[ID55f2eda4e05bd38d3320468960dcc37e]]
= Create or remove litigations for a path

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snaplock/litigations/{litigation.id}/operations`#

*Introduced In:* 9.7

Creates or removes litigations for the specified path.
```

```
== Required properties

* `type` - Legal-Hold operation type.
* `path` - Litigation path.

== Related ONTAP commands

* `snaplock legal-hold begin`
* `snaplock legal-hold end`

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_litigation.id_operations_
endpoint_overview.html[DOC
/storage/snaplock/litigations/{litigation.id}/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|litigation.id
|string
|path
|True
a|Litigation ID


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID.


|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.


|num_files_processed
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.


|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.
```

```
|type
|string
a|Specifies the type of legal-hold operation.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 30,
  "num_files_skipped": 10,
  "num_inodes_ignored": 10,
  "path": "/dir1",
  "state": "completed",
  "type": "begin"
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID.


|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.
```

```
|num_files_processed
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.


|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.


|type
|string
a|Specifies the type of legal-hold operation.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 30,
  "num_files_skipped": 10,
  "num_inodes_ignored": 10,
```

```
  "path": "/dir1",
  "state": "completed",
  "type": "begin"
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090343
| Invalid Field

| 14090346
| Internal Error. Wait a few minutes, then try the command again
|===


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#snaplock_legal_hold_operation]
[.api-collapsible-fifth-title]
snaplock_legal_hold_operation

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Operation ID.


|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.


|num_files_processed
```

```
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.


|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.


|type
|string
a|Specifies the type of legal-hold operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID89420ba9d3267cb87c4e185ca5ebb89e]]
= Cancel the ongoing legal-hold operation

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snaplock/litigations/{litigation.id}/operations/{id}`#

*Introduced In:* 9.7

Aborts the ongoing legal-hold operation. An abort does not rollback any
changes already made. You must re-run begin or end for cleanup.

== Related ONTAP commands

* `snaplock legal-hold abort`

== Example

----
DELETE "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-
0050568ebef5:l1/operations/16908292"
----

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_litigation.id_operations_
endpoint_overview.html[DOC
/storage/snaplock/litigations/{litigation.id}/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|litigation.id
|string
|path
```

```
|True
a|Litigation ID


|id
|string
|path
|True
a|Operation ID.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 13763378
| SnapLock legal-hold operation for this operation ID does not exist

| 14090346
| Internal Error. Wait a few minutes, then try the command again

| 14090541
| A completed or failed operation cannot be aborted
|===



[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDb38ede945e63a0cdb2dcc8c6df38ea98]]
= Retrieve the status of legal-hold for an operation ID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snaplock/litigations/{litigation.id}/operations/{id}`#

*Introduced In:* 9.7

Retrieves the status of legal-hold for the specified operation ID.

== Related ONTAP commands

* `snaplock legal-hold show`

== Learn more

*
xref:{relative_path}storage_snaplock_litigations_litigation.id_operations_
endpoint_overview.html[DOC
/storage/snaplock/litigations/{litigation.id}/operations]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|litigation.id
|string
|path
|True
a|Litigation ID


|id
|string
|path
|True
```

```
a|Operation ID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Operation ID.


|num_files_failed
|string
a|Specifies the number of files on which legal-hold operation failed.


|num_files_processed
|string
a|Specifies the number of files on which legal-hold operation was
successful.


|num_files_skipped
|string
a|Specifies the number of files on which legal-hold begin operation was
skipped. The legal-hold begin operation is skipped on a file if it is
already under hold for a given litigation.
```

```
|num_inodes_ignored
|string
a|Specifies the number of inodes on which the legal-hold operation was not
attempted because they were not regular files.


|path
|string
a|Specifies the path on which legal-hold operation is applied.


|state
|string
a|Specifies the status of legal-hold operation.


|type
|string
a|Specifies the type of legal-hold operation.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 16842759,
  "num_files_failed": 0,
  "num_files_processed": 30,
  "num_files_skipped": 10,
  "num_inodes_ignored": 10,
  "path": "/dir1",
  "state": "completed",
  "type": "begin"
}
====
```

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13763280
| Only a user with security login role "vsadmin-snaplock" is allowed to
perform this operation.

| 14090343
| Invalid Field

| 14090346
| Internal Error. Wait a few minutes, then try the command again
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


:leveloffset: -1


= SnapMirror

:leveloffset: +1


[[IDc2df80e23f5c68eebe19d48d59421b72]]
= ONTAP REST API SnapMirror endpoints



== Overview

SnapMirror technology enables customers to copy and backup their
production data. The secondary or destination volume in such a backup
relationship can reside anywhere, locally or remotely, and can be used to
restore access to the protected data. You can restore access to protected
data by drawing upon a vault of backups or, in the event that a primary is
unusable, by using a disaster recovery copy that can be activated for
clients and applications. SnapMirror periodically updates a replica to
create new backups and/or to keep a replica up-to-date with changes that
have been written to the primary. The SnapMirror subsystems are designed
to keep many pairs of source (primary) and destination (secondary) copies
up-to-date in an efficient and scalable manner.

The SnapMirror APIs can be used to create and manage SnapMirror
```

```
relationships of type "async", and "sync". These APIs can also be used to
manage restore-relationships. These APIs allow you to manage the following
endpoints:

* SnapMirror policies - When applied to a SnapMirror relationship, the
SnapMirror policy controls the behavior of the relationship and specifies
the configuration attributes for that relationship.
* SnapMirror relationships - You can create and manage SnapMirror
relationships, and you can change the state of the SnapMirror relationship
using a PATCH request.
* SnapMirror transfers - You can manage data transfers on the specified
SnapMirror relationship.





= Manage SnapMirror policies

:leveloffset: +1


[[IDb709d7765b6ae080627d27485fed5d9d]]
= Manage SnapMirror policies



== Managing SnapMirror policies

This API is used to manage SnapMirror policies of type "async" and "sync".
When applied to a SnapMirror relationship, the SnapMirror policy controls
the behavior of the relationship and specifies the configuration
attributes for that relationship.
The policy type "async" can be associated with a SnapMirror relationship
that has either a FlexVol volume, FlexGroup volume, or SVM as the
endpoint.
The policy type "sync" can be associated with a SnapMirror relationship
that has a FlexVol volume or a consistency group as the endpoint. The
policy type "sync" can have a "sync_type" of either "sync", "strict_sync",
"automated_failover" or "automated_failover_duplex". If the "sync_type" is
"sync", a write success is returned to the client after writing the data
to the primary endpoint and before writing the data to the secondary
endpoint. If the "sync_type" is "strict_sync", a write success is returned
to the client after writing the data to both primary and secondary
endpoints.
A "sync_type" of "automated_failover" can be associated with a SnapMirror
relationship that has a consistency group as the endpoint and provides
```

```
asymmetric active active access to the two storage copies.

A "sync_type" of "automated_failover_duplex" can be associated with a
SnapMirror relationship that has a consistency group as the endpoint and
provides symmetric active active access to the two storage copies.

Mapping of SnapMirror policies from CLI to REST

|===
| CLI | REST

| mirror-vault
| async

| async-mirror w/
| async w/

| all_source_snapshots
| copy_all_source_snapshots

| async-mirror w/o
| async w/

| all_source_snapshots
| copy_latest_source_snapshot

| vault
| async w/

|
| create_snapshot_on_source
|===

|                           |       |  sync_type                |
|                           |       |-------------------|
|sync-mirror                | sync  | sync                      |
|strict-sync-mirror         | sync  | strict_sync               |
|automated-failover         | sync  | automated_failover        |
|automated-failover-duplex  | sync  | automated_failover_duplex |
|                           |                                   |





[[IDfdf0c08fe4e115b66de344b25ad0c89a]]
= Retrieve SnapMirror async and sync policy types
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/policies`#

*Introduced In:* 9.6

Retrieves SnapMirror policies of type "async" and "sync".

== Related ONTAP commands

* `snapmirror policy show`

== Example

The following example shows how to retrieve a collection of SnapMirror
policies.

----
GET "/api/snapmirror/policies"
----

== Learn more

* xref:{relative_path}snapmirror_policies_endpoint_overview.html[DOC
/snapmirror/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|query
|False
a|Filter by uuid


|create_snapshot_on_source
|boolean
```

```
|query
|False
a|Filter by create_snapshot_on_source

* Introduced in: 9.11


|network_compression_enabled
|boolean
|query
|False
a|Filter by network_compression_enabled


|sync_type
|string
|query
|False
a|Filter by sync_type


|copy_latest_source_snapshot
|boolean
|query
|False
a|Filter by copy_latest_source_snapshot

* Introduced in: 9.11


|comment
|string
|query
|False
a|Filter by comment


|identity_preservation
|string
|query
|False
a|Filter by identity_preservation


|type
|string
|query
```

```
|False
a|Filter by type


|rpo
|integer
|query
|False
a|Filter by rpo

* Introduced in: 9.10


|name
|string
|query
|False
a|Filter by name


|scope
|string
|query
|False
a|Filter by scope


|copy_all_source_snapshots
|boolean
|query
|False
a|Filter by copy_all_source_snapshots

* Introduced in: 9.10


|transfer_schedule.name
|string
|query
|False
a|Filter by transfer_schedule.name


|transfer_schedule.uuid
|string
|query
|False
```

```
a|Filter by transfer_schedule.uuid


|throttle
|integer
|query
|False
a|Filter by throttle


|sync_common_snapshot_schedule.name
|string
|query
|False
a|Filter by sync_common_snapshot_schedule.name


|sync_common_snapshot_schedule.uuid
|string
|query
|False
a|Filter by sync_common_snapshot_schedule.uuid


|retention.count
|integer
|query
|False
a|Filter by retention.count


|retention.warn
|integer
|query
|False
a|Filter by retention.warn

* Introduced in: 9.13


|retention.period
|string
|query
|False
a|Filter by retention.period

* Introduced in: 9.11
```

```
|retention.creation_schedule.name
|string
|query
|False
a|Filter by retention.creation_schedule.name


|retention.creation_schedule.uuid
|string
|query
|False
a|Filter by retention.creation_schedule.uuid


|retention.preserve
|boolean
|query
|False
a|Filter by retention.preserve

* Introduced in: 9.13


|retention.label
|string
|query
|False
a|Filter by retention.label


|retention.prefix
|string
|query
|False
a|Filter by retention.prefix


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
```

```
|string
|query
|False
a|Filter by svm.uuid


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
```

```
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snapmirror_policy[snapmirror_policy]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
```

```
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "copy_all_source_snapshots": 1,
      "copy_latest_source_snapshot": 1,
      "create_snapshot_on_source": "",
      "identity_preservation": "string",
      "name": "Asynchronous",
      "retention": [
        {
          "count": 7,
          "creation_schedule": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "label": "hourly",
          "period": "P30D",
          "prefix": "string",
          "preserve": 1,
          "warn": 4
        }
      ],
      "rpo": 0,
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "sync_common_snapshot_schedule": {
        "_links": {
          "self": {
```

```
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "sync_type": "string",
      "throttle": 0,
      "transfer_schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "type": "string",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```


Status: Default


```
ONTAP Error Response codes

|===
| Error code | Description

| 13303842
| SnapMirror policy is not supported.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#creation_schedule]
[.api-collapsible-fifth-title]
creation_schedule

Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_policy_rule]
[.api-collapsible-fifth-title]
snapmirror_policy_rule

SnapMirror policy rule for retention.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of snapshots to be kept for retention. Maximum value will differ
based on type of relationship and scaling factor.


|creation_schedule
|link:#creation_schedule[creation_schedule]
a|Schedule used to create snapshots on the destination for long term
```

```
retention. Only cron schedules are supported for SnapMirror.


|label
|string
a|Snapshot label


|period
|string
a|Specifies the duration for which the snapshots are locked. The retention
period value represents a duration and must be in the ISO-8601 duration
format. Years, months, days, hours, minutes, and seconds are represented
as
"P+++<num>+++Y","P+++<num>+++M","P+++<num>+++D","PT+++<num>+++H","PT+++<nu
m>+++M" and "PT+++<num>+++S". Value "infinite" is also a valid input for
Flexvol volumes and FlexGroup volumes. A duration which combines different
periods is not supported, for example "P1Y10M" is not supported. The range
of supported retention period values is between 1 second to
infinite.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++</
num>+++


|prefix
|string
a|Specifies the prefix for the snapshot name to be created as per the
schedule. If no value is specified, then the label is used as the prefix.


|preserve
|boolean
a|Specifies the behavior when the snapshot retention count is reached on
the SnapMirror destination for the rule. The default value is false, which
means that the oldest snapshot is deleted to make room for new ones but
only if the number of snapshots has exceeded the retention count specified
in the 'count' property. When set to true and where the snapshots have
reached the retention count, an incremental SnapMirror transfer will fail
or if the rule has a schedule, snapshots will be no longer be created on
the SnapMirror destination.


|warn
|integer
a|Specifies the warning threshold count for the rule. The default value is
zero. When set to a value greater than zero, an event is generated after
the number of snapshots (for the particular rule) retained on a SnapMirror
destination reaches the specified warning limit. The preserve property for
```

```
the rule must be true in order to set the warn property to a value greater
than zero.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#sync_common_snapshot_schedule]
[.api-collapsible-fifth-title]
sync_common_snapshot_schedule

Schedule used to create common snapshots for synchronous relationships.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID
```

```
|===


[#snapmirror_policy]
[.api-collapsible-fifth-title]
snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type
"async", "sync" or "continuous".
The policy type "async" can be associated with a SnapMirror relationship
that has either a FlexVol volume, FlexGroup volume or SVM as the endpoint.
The policy type "sync" along with "sync_type" as "sync" or "strict_sync"
can be associated with a SnapMirror relationship that has a FlexVol volume
as the endpoint. The policy type "sync" can have a "sync_type" of either
"sync", "strict_sync", "automated_failover" or
"automated_failover_duplex". If the "sync_type" is "sync", a write success
is returned to the client after writing the data to the source endpoint
and before writing the data to the destination endpoint. If the
"sync_type" is "strict_sync", a write success is returned to the client
after writing the data to both source and destination endpoints.
If the "sync_type" is "automated_failover", the policy can be associated
with a SnapMirror active sync relationship that has a consistency group as
the endpoint and provides asymmetric active active access to the two
storage copies. If the "sync_type" is "automated_failover_duplex", the
policy can be associated with a SnapMirror active sync relationship that
has a consistency group as the endpoint and provides symmetric active
active access to the two storage copies. Use the "sync" policy with
"sync_type" as "automated_failover" or "automated_failover_duplex" to
establish SnapMirror active sync relationships for business continuity use
cases. SnapMirror relationships with policy types as "sync" and
"sync_type" as "automated_failover" or "automated_failover_duplex" can be
monitored by the Mediator, if configured. If the source Consistency Group
endpoint is not reachable, the Mediator might trigger a failover to the
destination consistency group endpoint.
A policy type of "continuous" can be associated with SnapMirror
relationships that have either ONTAP S3 buckets or non-ONTAP object stores
as endpoints. This type of policy is used for FabricLink owned targets.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|comment
|string
a|Comment associated with the policy.


|copy_all_source_snapshots
|boolean
a|Specifies that all the source snapshots (including the one created by
SnapMirror before the transfer begins) should be copied to the destination
on a transfer. "Retention" properties cannot be specified along with this
property. This is applicable only to async policies. Property can only be
set to 'true'.


|copy_latest_source_snapshot
|boolean
a|Specifies that the latest source snapshot (created by SnapMirror before
the transfer begins) should be copied to the destination on a transfer.
"Retention" properties cannot be specified along with this property. This
is applicable only to async policies. Property can only be set to 'true'.


|create_snapshot_on_source
|boolean
a|Specifies whether a new snapshot should be created on the source at the
beginning of an update or resync operation. This is applicable only to
async policies. Property can only be set to 'false'.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|name
|string
a|Name of the policy.


|network_compression_enabled
|boolean
```

```
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.


|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
a|Schedule used to create common snapshots for synchronous relationships.


|sync_type
|string
a|

|throttle
|integer
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.
```

```
|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID5757f7077155c99ca0f9e633b546ee54]]
= Create a SnapMirror policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/snapmirror/policies`#

*Introduced In:* 9.6

Creates a SnapMirror policy. The property "identity_preservation" is
applicable to only SnapMirror relationships with SVM endpoints and it
indicates which configuration of the source SVM is replicated to the
destination SVM.

It takes the following values:

* `full` - indicates that the source SVM configuration is replicated to
the destination SVM endpoint.
* `exclude_network_config` - indicates that the source SVM configuration
other than network configuration is replicated to the destination SVM
endpoint.
```

```
* `exclude_network_and_protocol_config` - indicates that the source SVM
configuration is not replicated to the destination SVM endpoint.

== Important notes

* The property "identity_preservation" is applicable to only SnapMirror
relationships with SVM endpoints and it indicates which configuration of
the source SVM is replicated to the destination SVM.
* The properties "identity_preservation" and "transfer_schedule" are not
applicable for "sync" type policies.
* The properties "copy_all_source_snapshots",
"copy_latest_source_snapshot", and "create_snapshot_on_source" are
mutually exclusive.
* The properties "copy_all_source_snapshots",
"copy_latest_source_snapshot", and "create_snapshot_on_source" are not
applicable for "sync" type policies.
* No "retention" properties can be specified if
"copy_all_source_snapshots" or 'copy_latest_source_snapshot' is specified.
* The properties "retention.creation_schedule" and "retention.prefix" are
not applicable for "sync" type policies.
* The property "retention.creation_schedule" is not applicable for "async"
policies with "create_snapshot_on_source" set to "false".
* The property "sync_common_snapshot_schedule" is not applicable for an
"async" type policy.
* The property "retention.count" specifies the maximum number of snapshots
that are retained on the SnapMirror destination volume.
* When the property "retention.label" is specified, the snapshots that
have a SnapMirror label matching this property is transferred to the
SnapMirror destination.
* When the property "retention.creation_schedule" is specified, snapshots
are directly created on the SnapMirror destination. The snapshots created
have the same content as the latest snapshot already present on the
SnapMirror destination.
* The property "transfer_schedule" cannot be set to null (no-quotes)
during SnapMirror policy POST.
* The properties "retention.label" and "retention.count" must be specified
for "async" policies with "create_snapshot_on_source" set to "false".
* The property "retention.warn" is not supported for a policy when the
"retention.preserve" property is false.
* The property "retention.warn" value must be less than the property
"retention.count" value for a rule in a policy.

== Required properties

* `name` - Name of the new SnapMirror policy.
```

```
== Recommended optional properties

* `svm.name` or `svm.uuid` - Name or UUID of the SVM that owns the
SnapMirror policy.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `type` - _async_
* `sync_type` - _sync_ (when `type` is _sync_)
* `network_compression_enabled` - _false_
* `throttle` - _0_
* `identity_preservation` - `_exclude_network_and_protocol_config_`

== Related ONTAP commands

* `snapmirror policy create`

== Examples

Creating a SnapMirror policy of type "sync"

----
 POST "/api/snapmirror/policies/" '{"name": "policy1", "svm.name": "VS0",
"type": "sync", "sync_type": "sync"}'
----

Creating a SnapMirror policy of type "async" with two sets of retention
values, one with a creation_schedule

----
 POST "/api/snapmirror/policies" '{"name": "policy_ret", "svm": {"name":
"vs1"}, "retention": [{"label": "weekly", "count": "2",
"creation_schedule": {"name": "weekly"}}, {"label":"daily",
"count":"7"}]}'
----

Creating a SnapMirror policy of type "async"

----
 POST "/api/snapmirror/policies" '{"name": "newPolicy", "svm":{"name" :
"vs1"}, "type": "async"}'
----

Creating a SnapMirror policy of type "async" which replicates all
```

```
snapshots

----
 POST "/api/snapmirror/policies" '{"name": "newPolicy", "svm":{"name" :
"vs1"}, "type": "async", "copy_all_source_snapshots": "true"}'
----

Creating a SnapMirror policy of type "async" which replicates latest
snapshot

----
 POST "/api/snapmirror/policies" '{"name": "newPolicy2", "svm":{"name" :
"vs1"}, "type": "async", "copy_latest_source_snapshot": "true"}'
----

Creating a SnapMirror policy of type "async" which does not create
snapshots on source

----
 POST "/api/snapmirror/policies" '{"name": "newPolicy", "svm":{"name" :
"vs1"}, "type": "async", "create_snapshot_on_source": "false",
"retention": [{"label": "daily", "count": 7}]}'
----

Creating a SnapMirror policy of type "sync" with sync_type as
"automated_failover"

----
 POST "/api/snapmirror/policies/" '{"name": "policy1", "svm.name": "VS0",
"type": "sync", "sync_type": "automated_failover" }'
----

Creating a SnapMirror policy of type "sync" with sync_type as
"automated_failover_duplex"

----
 POST "/api/snapmirror/policies/" '{"name": "policy_afd", "svm.name":
"VS0", "type": "sync", "sync_type": "automated_failover_duplex" }'
----

Creating a SnapMirror policy of type "async" with two sets of retention
values and retention periods

----
 POST "/api/snapmirror/policies" '{"name": "policy_ret", "svm": {"name":
"vs1"}, "retention": [{"label": "weekly", "count": "2", "period": "P7D"},
{"label":"daily", "count":"7", "period": "PT3H"}]}'
```

```
----

Creating a SnapMirror policy of type "async" with retention value as
"infinite"

----
 POST "/api/snapmirror/policies" '{"name": "policy_ret", "svm": {"name":
"vs1"}, "retention": [{"label": "weekly", "count": "5", "period":
"infinite"}]}'
----

Creating a SnapMirror policy of type "async" with properties retention
preserve as true and retention warn as 3.

----
 POST "/api/snapmirror/policies" '{"name": "policy_ret", "svm": {"name":
"vs1"}, "retention": [{"label": "weekly", "count": "5", "preserve": true,
"warn": 3}]}'
----

== Learn more

* xref:{relative_path}snapmirror_policies_endpoint_overview.html[DOC
/snapmirror/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:


|return_timeout
```

```
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment associated with the policy.


|copy_all_source_snapshots
|boolean
a|Specifies that all the source snapshots (including the one created by
SnapMirror before the transfer begins) should be copied to the destination
on a transfer. "Retention" properties cannot be specified along with this
property. This is applicable only to async policies. Property can only be
set to 'true'.


|copy_latest_source_snapshot
|boolean
a|Specifies that the latest source snapshot (created by SnapMirror before
the transfer begins) should be copied to the destination on a transfer.
"Retention" properties cannot be specified along with this property. This
is applicable only to async policies. Property can only be set to 'true'.
```

```
|create_snapshot_on_source
|boolean
a|Specifies whether a new snapshot should be created on the source at the
beginning of an update or resync operation. This is applicable only to
async policies. Property can only be set to 'false'.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|name
|string
a|Name of the policy.


|network_compression_enabled
|boolean
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.


|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".


|svm
```

```
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
a|Schedule used to create common snapshots for synchronous relationships.


|sync_type
|string
a|

|throttle
|integer
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "copy_all_source_snapshots": 1,
  "copy_latest_source_snapshot": 1,
  "create_snapshot_on_source": "",
  "identity_preservation": "string",
  "name": "Asynchronous",
```

```
  "retention": [
    {
      "count": 7,
      "creation_schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "period": "P30D",
      "prefix": "string",
      "preserve": 1,
      "warn": 4
    }
  ],
  "rpo": 0,
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "sync_common_snapshot_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "sync_type": "string",
  "throttle": 0,
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```


Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 6619714
| Schedule specified is an interval schedule. SnapMirror does not support
interval schedules.

| 13304126
| Enter a value for "count" in the range indicated in the error message.

| 13303887
| Failed to create SnapMirror policy. Reason: Maximum number of allowed
retention rules reached

| 13304083
| The specified property is not supported because all nodes in the cluster
are not capable of supporting this property.

| 13304084
| Properties specified are mutually exclusive. Provide only one property.

| 13304085
| The specified property does not support the specified value.

| 13304092
| Input value of the retention period property is invalid. For
relationships with FlexVol volume or FlexGroup volume destinations, the
duration must be in ISO 6801 format or can be infinite. For relationships
with object store destinations, only duration values with Y, M or D and
supported and must be in the specified range.
|===

| 6621045     | The property retention.warn is not supported for a policy
when the retention.preserve property is false.
| 13304129    | The property retention.warn value must be less than the
property retention.count value for a rule in a policy.
| 13304130    | The total retention.count value for all rules in a policy
cannot exceed the value indicated in the error message.
| 6621060     | Failed to create specified policy. Reason: A policy with
```

```
this name already exists. |
| 13304118    | The specified properties cannot be both set at the same
time:  retention.creation_schedule and retention.period |
| 6621972     | Schedule specified is not supported by SnapMirror
Synchronous. The allowed time interval for the creation of common
snapshots can range between 30 minutes and 24 hours. |
| 13304011    | Failed to create or modify the specified SnapMirror
policy. Reason: The property retention.count cannot have a value greater
than 1. |
| 6621091     | The specified SnapMirror policy cannot have a rule with a
preserve value of true. |
| 6621088     | The specified SnapMirror policy must have at least one
rule without a schedule. |


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#creation_schedule]
[.api-collapsible-fifth-title]
creation_schedule
```

```
Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_policy_rule]
[.api-collapsible-fifth-title]
snapmirror_policy_rule

SnapMirror policy rule for retention.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of snapshots to be kept for retention. Maximum value will differ
based on type of relationship and scaling factor.


|creation_schedule
|link:#creation_schedule[creation_schedule]
a|Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.
```

```
|label
|string
a|Snapshot label


|period
|string
a|Specifies the duration for which the snapshots are locked. The retention
period value represents a duration and must be in the ISO-8601 duration
format. Years, months, days, hours, minutes, and seconds are represented
as
"P+++<num>+++Y","P+++<num>+++M","P+++<num>+++D","PT+++<num>+++H","PT+++<nu
m>+++M" and "PT+++<num>+++S". Value "infinite" is also a valid input for
Flexvol volumes and FlexGroup volumes. A duration which combines different
periods is not supported, for example "P1Y10M" is not supported. The range
of supported retention period values is between 1 second to
infinite.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++</
num>+++


|prefix
|string
a|Specifies the prefix for the snapshot name to be created as per the
schedule. If no value is specified, then the label is used as the prefix.


|preserve
|boolean
a|Specifies the behavior when the snapshot retention count is reached on
the SnapMirror destination for the rule. The default value is false, which
means that the oldest snapshot is deleted to make room for new ones but
only if the number of snapshots has exceeded the retention count specified
in the 'count' property. When set to true and where the snapshots have
reached the retention count, an incremental SnapMirror transfer will fail
or if the rule has a schedule, snapshots will be no longer be created on
the SnapMirror destination.


|warn
|integer
a|Specifies the warning threshold count for the rule. The default value is
zero. When set to a value greater than zero, an event is generated after
the number of snapshots (for the particular rule) retained on a SnapMirror
destination reaches the specified warning limit. The preserve property for
the rule must be true in order to set the warn property to a value greater
```

```
than zero.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#sync_common_snapshot_schedule]
[.api-collapsible-fifth-title]
sync_common_snapshot_schedule

Schedule used to create common snapshots for synchronous relationships.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_policy]
[.api-collapsible-fifth-title]
snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type
"async", "sync" or "continuous".
The policy type "async" can be associated with a SnapMirror relationship
that has either a FlexVol volume, FlexGroup volume or SVM as the endpoint.
```

```
The policy type "sync" along with "sync_type" as "sync" or "strict_sync"
can be associated with a SnapMirror relationship that has a FlexVol volume
as the endpoint. The policy type "sync" can have a "sync_type" of either
"sync", "strict_sync", "automated_failover" or
"automated_failover_duplex". If the "sync_type" is "sync", a write success
is returned to the client after writing the data to the source endpoint
and before writing the data to the destination endpoint. If the
"sync_type" is "strict_sync", a write success is returned to the client
after writing the data to both source and destination endpoints.
If the "sync_type" is "automated_failover", the policy can be associated
with a SnapMirror active sync relationship that has a consistency group as
the endpoint and provides asymmetric active active access to the two
storage copies. If the "sync_type" is "automated_failover_duplex", the
policy can be associated with a SnapMirror active sync relationship that
has a consistency group as the endpoint and provides symmetric active
active access to the two storage copies. Use the "sync" policy with
"sync_type" as "automated_failover" or "automated_failover_duplex" to
establish SnapMirror active sync relationships for business continuity use
cases. SnapMirror relationships with policy types as "sync" and
"sync_type" as "automated_failover" or "automated_failover_duplex" can be
monitored by the Mediator, if configured. If the source Consistency Group
endpoint is not reachable, the Mediator might trigger a failover to the
destination consistency group endpoint.
A policy type of "continuous" can be associated with SnapMirror
relationships that have either ONTAP S3 buckets or non-ONTAP object stores
as endpoints. This type of policy is used for FabricLink owned targets.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment associated with the policy.


|copy_all_source_snapshots
|boolean
a|Specifies that all the source snapshots (including the one created by
SnapMirror before the transfer begins) should be copied to the destination
on a transfer. "Retention" properties cannot be specified along with this
property. This is applicable only to async policies. Property can only be
set to 'true'.
```

```
|copy_latest_source_snapshot
|boolean
a|Specifies that the latest source snapshot (created by SnapMirror before
the transfer begins) should be copied to the destination on a transfer.
"Retention" properties cannot be specified along with this property. This
is applicable only to async policies. Property can only be set to 'true'.


|create_snapshot_on_source
|boolean
a|Specifies whether a new snapshot should be created on the source at the
beginning of an update or resync operation. This is applicable only to
async policies. Property can only be set to 'false'.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|name
|string
a|Name of the policy.


|network_compression_enabled
|boolean
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.
```

```
|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
a|Schedule used to create common snapshots for synchronous relationships.


|sync_type
|string
a|

|throttle
|integer
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID44602ad5dc3fdfd5844bf1cca0c95a1a]]
= Delete a SnapMirror policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/snapmirror/policies/{uuid}`#

*Introduced In:* 9.6

Deletes a SnapMirror policy.

== Related ONTAP commands

* `snapmirror policy delete`

== Example

----
DELETE "/api/snapmirror/policies/510c15d4-f9e6-11e8-bdb5-0050568e12c2"
```

```
----

== Learn more

* xref:{relative_path}snapmirror_policies_endpoint_overview.html[DOC
/snapmirror/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|SnapMirror policy UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDcc63bca3c5233117bc5d719f34f4fcaa]]
= Retrieve a specific SnapMirror policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/policies/{uuid}`#

*Introduced In:* 9.6

Retrieves a specific SnapMirror policy.

== Example

----
GET "/api/snapmirror/policies/567aaac0-f863-11e8-a666-0050568e12c2"
----

== Learn more

* xref:{relative_path}snapmirror_policies_endpoint_overview.html[DOC
/snapmirror/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
```

```
|path
|True
a|SnapMirror policy UUID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|Comment associated with the policy.


|copy_all_source_snapshots
|boolean
a|Specifies that all the source snapshots (including the one created by
SnapMirror before the transfer begins) should be copied to the destination
on a transfer. "Retention" properties cannot be specified along with this
property. This is applicable only to async policies. Property can only be
set to 'true'.


|copy_latest_source_snapshot
|boolean
a|Specifies that the latest source snapshot (created by SnapMirror before
the transfer begins) should be copied to the destination on a transfer.
"Retention" properties cannot be specified along with this property. This
```

```
is applicable only to async policies. Property can only be set to 'true'.


|create_snapshot_on_source
|boolean
a|Specifies whether a new snapshot should be created on the source at the
beginning of an update or resync operation. This is applicable only to
async policies. Property can only be set to 'false'.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|name
|string
a|Name of the policy.


|network_compression_enabled
|boolean
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.


|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
a|Schedule used to create common snapshots for synchronous relationships.


|sync_type
|string
a|

|throttle
|integer
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```
  "comment": "string",
  "copy_all_source_snapshots": 1,
  "copy_latest_source_snapshot": 1,
  "create_snapshot_on_source": "",
  "identity_preservation": "string",
  "name": "Asynchronous",
  "retention": [
    {
      "count": 7,
      "creation_schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "period": "P30D",
      "prefix": "string",
      "preserve": 1,
      "warn": 4
    }
  ],
  "rpo": 0,
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "sync_common_snapshot_schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "sync_type": "string",
```

```
  "throttle": 0,
  "transfer_schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303842
| SnapMirror policy is not supported.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#creation_schedule]
[.api-collapsible-fifth-title]
creation_schedule

Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_policy_rule]
[.api-collapsible-fifth-title]
snapmirror_policy_rule

SnapMirror policy rule for retention.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of snapshots to be kept for retention. Maximum value will differ
based on type of relationship and scaling factor.


|creation_schedule
|link:#creation_schedule[creation_schedule]
a|Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.


|label
|string
a|Snapshot label


|period
|string
a|Specifies the duration for which the snapshots are locked. The retention
period value represents a duration and must be in the ISO-8601 duration
format. Years, months, days, hours, minutes, and seconds are represented
as
"P+++<num>+++Y","P+++<num>+++M","P+++<num>+++D","PT+++<num>+++H","PT+++<nu
m>+++M" and "PT+++<num>+++S". Value "infinite" is also a valid input for
Flexvol volumes and FlexGroup volumes. A duration which combines different
periods is not supported, for example "P1Y10M" is not supported. The range
of supported retention period values is between 1 second to
infinite.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++</
num>+++


|prefix
|string
a|Specifies the prefix for the snapshot name to be created as per the
schedule. If no value is specified, then the label is used as the prefix.


|preserve
```

```
|boolean
a|Specifies the behavior when the snapshot retention count is reached on
the SnapMirror destination for the rule. The default value is false, which
means that the oldest snapshot is deleted to make room for new ones but
only if the number of snapshots has exceeded the retention count specified
in the 'count' property. When set to true and where the snapshots have
reached the retention count, an incremental SnapMirror transfer will fail
or if the rule has a schedule, snapshots will be no longer be created on
the SnapMirror destination.


|warn
|integer
a|Specifies the warning threshold count for the rule. The default value is
zero. When set to a value greater than zero, an event is generated after
the number of snapshots (for the particular rule) retained on a SnapMirror
destination reaches the specified warning limit. The preserve property for
the rule must be true in order to set the warn property to a value greater
than zero.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
```

```
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#sync_common_snapshot_schedule]
[.api-collapsible-fifth-title]
sync_common_snapshot_schedule

Schedule used to create common snapshots for synchronous relationships.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID92eba8740efb8d7e697b509a3d70c88c]]
= Update the SnapMirror policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/snapmirror/policies/{uuid}`#

*Introduced In:* 9.6

Updates the SnapMirror policy.

== Important notes
```

```
* The properties "retention.label" and "retention.count" are mandatory if
"retention" is provided in the input. The provided "retention.label" is
the final list and it replaces the existing values.
* The value of the "identity_preservation" property cannot be changed if
the SnapMirror relationships associated with the policy have different
identity_preservation configurations.
* If the SnapMirror policy "identity_preservation" value matches the
"identity_preservation" value of the associated SnapMirror relationships,
then the "identity_preservation" value can be changed from a higher
"identity_preservation" threshold value to a lower "identity_preservation"
threshold value but not vice-versa. For example, the threshold value of
the "identity_preservation" property can be changed from "full" to
"exclude_network_config", but cannot be increased from
"exclude_network_and_protocol_config" to "exclude_network_config" to
"full". The threshold value of the "identity_preservation" cannot be
changed to "exclude_network_and_protocol_config" for IDP SVMDR.
* The policy properties "copy_all_source_snapshots",
"copy_latest_source_snapshot", and "create_snapshot_on_source" cannot be
modified.
* No "retention" properties can be modified if the
"copy_all_source_snapshots" or "copy_latest_source_snapshot" property is
present in the policy.
* Replacing or deleting all retention rules of a policy that has the
"create_snapshot_on_source" property set to false in a single API call is
not supported.
* Modifying the property "retention.label" for all retention rules of a
policy that has the "create_snapshot_on_source" property set to false in a
single API call is not supported.
* To remove a transfer_schedule on a SnapMirror policy set the
"transfer_schedule" to null (no-quotes) during SnapMirror policy PATCH.

== Related ONTAP commands

* `snapmirror policy modify`

== Example

Updating the "retention" property to add rules to a policy without any
rules.

----
 PATCH "/api/snapmirror/policies/fe65686d-00dc-11e9-b5fb-0050568e3f83"
'{"retention": [{"label": "newlabel", "count": 2}, {"label": "weekly",
"count": 2, "creation_schedule": {"name": "weekly"}}, {"label": "daily",
"count": 14}]}'
```

```
----

Updating the "retention" property to add rules to a policy with existing
rules {"retention": [{"label": "oldLabel1", "count": 2}, {"label":
"oldLabel2", "count": 5}]

----
 PATCH "/api/snapmirror/policies/fe65686d-00dc-11e9-b5fb-0050568e3f83"
'{"retention": [{"label": "oldLabel1", "count": 2}, {"label": "oldLabel2",
"count": 5}, {"label": "newlabel", "count": 3}, {"label": "weekly",
"count": 1}]}'
----

Updating the "retention" property to remove a rule (oldLabel1) and add new
rule to a policy with existing rules {"retention": [{"label": "oldLabel1",
"count": 2}, {"label": "oldLabel2", "count": 3}]

----
 PATCH "/api/snapmirror/policies/fe65686d-00dc-11e9-b5fb-0050568e3f83"
'{"retention": [{"label": "oldLabel2", "count": 3}, {"label": "newlabel",
"count": 2}]}'
----

Updating "transfer_schedule", "throttle", and "identity_preservation"
properties

----
 PATCH "/api/snapmirror/policies/8aef950b-3bef-11e9-80ac-0050568ea591"
'{"transfer_schedule.name" : "weekly", "throttle" : "100",
"identity_preservation":"exclude_network_and_protocol_config"}'
----

Removing the SnapMirror transfer_schedule for a SnapMirror policy.
Transfer_schedule can be specified as UUID or name or both with the value
set to null (no-quotes).

----
 PATCH "/api/snapmirror/policies/98bb2608-fc60-11e8-aa13-005056a707ff/"
'{"transfer_schedule":{"uuid":null, "name":null}}'
----

Updating the "retention" property to have retention.preserve and
retention.warn for existing rule.

----
 PATCH "/api/snapmirror/policies/fe65686d-00dc-11e9-b5fb-0050568e3f83"
'{"retention": [{"label": "oldLabel1", "count": 3, "preserve": true,
```

```
"warn": 2}]}'
----

== Learn more

* xref:{relative_path}snapmirror_policies_endpoint_overview.html[DOC
/snapmirror/policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|SnapMirror policy UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Comment associated with the policy.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|network_compression_enabled
|boolean
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.


|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
```

```
a|Schedule used to create common snapshots for synchronous relationships.


|throttle
|integer
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "copy_all_source_snapshots": 1,
  "copy_latest_source_snapshot": 1,
  "create_snapshot_on_source": "",
  "identity_preservation": "string",
  "retention": [
    {
      "count": 7,
      "creation_schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "period": "P30D",
      "prefix": "string",
      "preserve": 1,
      "warn": 4
    }
  ],
  "rpo": 0,
```

```
  "scope": "string",
  "sync_common_snapshot_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "throttle": 0,
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 6619714
| Schedule specified is an interval schedule. SnapMirror does not support
interval schedules.

| 13303842
| SnapMirror policy is not supported.

| 13303843
| Conflicting values between SnapMirror policy and SnapMirror
relationships for either 'transfer_schedule, throttle or
identity_preservation' properties

| 13303850
| Invalid input parameter

| 13303887
| Failed to create SnapMirror policy. Reason: Maximum number of allowed
retention rules reached

| 13304050
| Retention cannot be empty for a SnapMirror policy with
'create_snapshot_on_source' set to false.

| 13304092
| Input value of the retention period property is invalid. For
relationships with FlexVol volume or FlexGroup volume destinations, the
duration must be in ISO 6801 format or can be infinite. For relationships
with object store destinations, only duration values with Y, M or D and
supported and must be in the specified range.

| 6621045
| The property retention.warn is not supported for a policy when the
retention.preserve property is false.

| 13304129
| The property retention.warn value must be less than the property
```

```
retention.count value for a rule in a policy.

| 13304126
| Enter a value for "count" in the range indicated in the error message.

| 13304130
| The total retention.count value for all rules in a policy cannot exceed
the value indicated in the error message."

| 13304131
| Modifying property "retention.creation_schedule" for the policy is not
supported because the policy is associated with a relationship with a
FlexGroup volume endpoint.

| 13304011
| Failed to create or modify the specified Snapmirror policy. Reason: The
property retention.count cannot have a value greater than 1.

| 6621091
| The specified Snapmirror policy cannot have a rule with a preserve value
of true.

| 6621088
| The specified Snapmirror policy must have at least one rule without a
schedule.

| 6621129
| Cannot add a snapshot creation schedule to the policy because it is
associated with a Consistency Group relationship.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#creation_schedule]
[.api-collapsible-fifth-title]
creation_schedule

Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_policy_rule]
[.api-collapsible-fifth-title]
snapmirror_policy_rule

SnapMirror policy rule for retention.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|count
|integer
a|Number of snapshots to be kept for retention. Maximum value will differ
based on type of relationship and scaling factor.


|creation_schedule
|link:#creation_schedule[creation_schedule]
a|Schedule used to create snapshots on the destination for long term
retention. Only cron schedules are supported for SnapMirror.


|label
|string
a|Snapshot label


|period
|string
a|Specifies the duration for which the snapshots are locked. The retention
period value represents a duration and must be in the ISO-8601 duration
format. Years, months, days, hours, minutes, and seconds are represented
as
"P+++<num>+++Y","P+++<num>+++M","P+++<num>+++D","PT+++<num>+++H","PT+++<nu
m>+++M" and "PT+++<num>+++S". Value "infinite" is also a valid input for
Flexvol volumes and FlexGroup volumes. A duration which combines different
periods is not supported, for example "P1Y10M" is not supported. The range
of supported retention period values is between 1 second to
infinite.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>++++++</
num>+++


|prefix
|string
a|Specifies the prefix for the snapshot name to be created as per the
schedule. If no value is specified, then the label is used as the prefix.


|preserve
|boolean
a|Specifies the behavior when the snapshot retention count is reached on
the SnapMirror destination for the rule. The default value is false, which
```

```
means that the oldest snapshot is deleted to make room for new ones but
only if the number of snapshots has exceeded the retention count specified
in the 'count' property. When set to true and where the snapshots have
reached the retention count, an incremental SnapMirror transfer will fail
or if the rule has a schedule, snapshots will be no longer be created on
the SnapMirror destination.


|warn
|integer
a|Specifies the warning threshold count for the rule. The default value is
zero. When set to a value greater than zero, an event is generated after
the number of snapshots (for the particular rule) retained on a SnapMirror
destination reaches the specified warning limit. The preserve property for
the rule must be true in order to set the warn property to a value greater
than zero.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#sync_common_snapshot_schedule]
[.api-collapsible-fifth-title]
sync_common_snapshot_schedule

Schedule used to create common snapshots for synchronous relationships.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
```

```
|string
a|Job schedule UUID


|===


[#snapmirror_policy]
[.api-collapsible-fifth-title]
snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type
"async", "sync" or "continuous".
The policy type "async" can be associated with a SnapMirror relationship
that has either a FlexVol volume, FlexGroup volume or SVM as the endpoint.
The policy type "sync" along with "sync_type" as "sync" or "strict_sync"
can be associated with a SnapMirror relationship that has a FlexVol volume
as the endpoint. The policy type "sync" can have a "sync_type" of either
"sync", "strict_sync", "automated_failover" or
"automated_failover_duplex". If the "sync_type" is "sync", a write success
is returned to the client after writing the data to the source endpoint
and before writing the data to the destination endpoint. If the
"sync_type" is "strict_sync", a write success is returned to the client
after writing the data to both source and destination endpoints.
If the "sync_type" is "automated_failover", the policy can be associated
with a SnapMirror active sync relationship that has a consistency group as
the endpoint and provides asymmetric active active access to the two
storage copies. If the "sync_type" is "automated_failover_duplex", the
policy can be associated with a SnapMirror active sync relationship that
has a consistency group as the endpoint and provides symmetric active
active access to the two storage copies. Use the "sync" policy with
"sync_type" as "automated_failover" or "automated_failover_duplex" to
establish SnapMirror active sync relationships for business continuity use
cases. SnapMirror relationships with policy types as "sync" and
"sync_type" as "automated_failover" or "automated_failover_duplex" can be
monitored by the Mediator, if configured. If the source Consistency Group
endpoint is not reachable, the Mediator might trigger a failover to the
destination consistency group endpoint.
A policy type of "continuous" can be associated with SnapMirror
relationships that have either ONTAP S3 buckets or non-ONTAP object stores
as endpoints. This type of policy is used for FabricLink owned targets.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|comment
|string
a|Comment associated with the policy.


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type.


|network_compression_enabled
|boolean
a|Specifies whether network compression is enabled for transfers. This is
applicable only to the policies of type "async".


|retention
|array[link:#snapmirror_policy_rule[snapmirror_policy_rule]]
a|Rules for snapshot retention.


|rpo
|integer
a|Specifies the duration of time for which a change to be propagated to a
mirror should be delayed, in seconds. This is an intentional propagation
delay between mirrors and is configurable down to zero, which means an
immediate propagation. This is supported for policies of type
'continuous'.


|scope
|string
a|Set to "svm" for policies owned by an SVM, otherwise set to "cluster".


|sync_common_snapshot_schedule
|link:#sync_common_snapshot_schedule[sync_common_snapshot_schedule]
a|Schedule used to create common snapshots for synchronous relationships.


|throttle
|integer
```

```
a|Throttle in KB/s. Default to unlimited.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|The schedule used to update asynchronous relationships. Only cron
schedules are supported for SnapMirror.


|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
:leveloffset: -1


= Manage SnapMirror relationships

:leveloffset: +1


[[ID332500ee9307ff5042f4d2c0858a1f6f]]
= Manage SnapMirror relationships



== Overview

This API manages asynchronous extended data protection (XDP) relationships
for FlexVol volumes, FlexGroup volumes, or SVMs. It is also used to manage
a synchronous relationship between FlexVol volumes, which provides zero
RPO data protection and active synchronous relationship with automated
failover between Application Consistency Group endpoints which provides
zero RTO data protection.

To create an asynchronous extended data protection relationship with
FlexVol volumes, FlexGroup volumes, Application Consistency Groups or
SVMs, use the policy of type "async".

To create a synchronous relationship between FlexVol volumes, use the
policy of type "sync" with sync_type of either "sync" or "strict_sync". To
create an active synchronous relationship with automated failover between
Application Consistency Group endpoints, use the policy of type "sync"
with sync_type "automated_failover". You can create an asynchronous
extended data protection relationship between the source and destination
which can be used by the transfer APIs to perform SnapMirror "restore"
operations.

To create FlexVol volume or FlexGroup volume SnapMirror relationships, the
source volume must be in the "online" state and be a read_write type; the
destination volume must be in the "online" state and be a data protection
type.

In the case of an asynchronous or synchronous SnapMirror relationship for
an Application Consistency Group of FlexVol volumes, SnapMirror creation
results in the creation of an Application Consistency Group on the source
cluster if it did not already exist with the exact same name and set of
```

```
FlexVol volumes specified in the current operation. Additionally, if the
specified Application Consistency Group is already present and is already
a part of an existing SnapMirror relationship, the process fails. Creating
an Application Consistency Group on the destination cluster is part of the
SnapMirror creation workflow.

To create SnapMirror relationships between SVMs, the source SVM must be of
subtype "default" and the destination SVM of subtype "dp_destination".
Additionally, SVMs must be peered before a relationship can be established
between them when the "create_destination" property is not specified. When
the "create_destination" property is specified, the destination SVM is
provisioned on the destination cluster and the SVM peer relationship is
established between the source SVM and the new destination SVM, provided
that the source SVM has SVM peering permissions for the destination
cluster.

Data protection FlexVol volume SnapMirror relationships cannot be created
using this API but existing relationships can be listed or managed.

The SnapMirror functionality is subdivided into relationship APIs and
transfer APIs:

* SnapMirror relationship APIs are used to create and manage the
SnapMirror relationships.
* SnapMirror transfer APIs are used to manage data transfers.




[[ID829034a6acb477a5610b17b33021ba47]]
= Retrieve information for SnapMirror relationships

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/relationships`#

*Introduced In:* 9.6

Retrieves information for SnapMirror relationships whose destination
endpoints are in the current SVM or the current cluster, depending on the
cluster context.

== Related ONTAP commands

* `snapmirror show`
* `snapmirror list-destinations`
```

```
== Expensive properties

* `source.consistency_group_volumes.name`
* `destination.consistency_group_volumes.name`
* `svmdr_volumes.name`

== Examples

The following examples show how to retrieve the list of SnapMirror
relationships and the list of SnapMirror destinations.

. Retrieving the list of SnapMirror relationships. This API must be run on
the cluster containing the destination endpoint.

----
 GET "/api/snapmirror/relationships/"
----

. Retrieving the list of SnapMirror destinations on source. This must be
run on the cluster containing the source endpoint.

----
 GET "/api/snapmirror/relationships/?list_destinations_only=true"
----

. Retrieving the relationship UUID of SnapMirror relationships with lag
time greater than 2 days. This API must be run on the cluster containing
the destination endpoint.

----
 GET "/api/snapmirror/relationships/?fields=uuid&lag_time=>P2DT"
----

. Retrieving the list of SnapMirror relationships with lag time less than
10 hours. This API must be run on the cluster containing the destination
endpoint.

----
 GET "/api/snapmirror/relationships/?lag_time=<PT10H"
----

. Retrieving the list of constituent volumes for SVM DR Snapmirror
relationships.

----
 GET "/api/snapmirror/relationships/?fields=svmdr_volumes.name"
----
```

```
</private>

== Learn more

* xref:{relative_path}snapmirror_relationships_endpoint_overview.html[DOC
/snapmirror/relationships]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|list_destinations_only
|boolean
|query
|False
a|Set to true to show relationships from the source only.


|transfer.end_time
|string
|query
|False
a|Filter by transfer.end_time

* Introduced in: 9.9


|transfer.uuid
|string
|query
|False
a|Filter by transfer.uuid


|transfer.bytes_transferred
|integer
|query
|False
```

```
a|Filter by transfer.bytes_transferred


|transfer.state
|string
|query
|False
a|Filter by transfer.state


|transfer.total_duration
|string
|query
|False
a|Filter by transfer.total_duration

* Introduced in: 9.9


|transfer.type
|string
|query
|False
a|Filter by transfer.type

* Introduced in: 9.14


|transfer.last_updated_time
|string
|query
|False
a|Filter by transfer.last_updated_time

* Introduced in: 9.14


|state
|string
|query
|False
a|Filter by state


|restore
|boolean
|query
```

```
|False
a|Filter by restore


|policy.name
|string
|query
|False
a|Filter by policy.name


|policy.type
|string
|query
|False
a|Filter by policy.type


|policy.uuid
|string
|query
|False
a|Filter by policy.uuid


|source.cluster.name
|string
|query
|False
a|Filter by source.cluster.name

* Introduced in: 9.7


|source.cluster.uuid
|string
|query
|False
a|Filter by source.cluster.uuid

* Introduced in: 9.7


|source.consistency_group_volumes.name
|string
|query
|False
```

```
a|Filter by source.consistency_group_volumes.name

* Introduced in: 9.8


|source.path
|string
|query
|False
a|Filter by source.path


|source.luns.uuid
|string
|query
|False
a|Filter by source.luns.uuid

* Introduced in: 9.16


|source.luns.name
|string
|query
|False
a|Filter by source.luns.name

* Introduced in: 9.16


|source.svm.name
|string
|query
|False
a|Filter by source.svm.name


|source.svm.uuid
|string
|query
|False
a|Filter by source.svm.uuid


|lag_time
|string
|query
```

```
|False
a|Filter by lag_time


|consistency_group_failover.error.arguments.message
|string
|query
|False
a|Filter by consistency_group_failover.error.arguments.message

* Introduced in: 9.8


|consistency_group_failover.error.arguments.code
|string
|query
|False
a|Filter by consistency_group_failover.error.arguments.code

* Introduced in: 9.8


|consistency_group_failover.error.code
|string
|query
|False
a|Filter by consistency_group_failover.error.code

* Introduced in: 9.8


|consistency_group_failover.error.message
|string
|query
|False
a|Filter by consistency_group_failover.error.message

* Introduced in: 9.8


|consistency_group_failover.state
|string
|query
|False
a|Filter by consistency_group_failover.state

* Introduced in: 9.14
```

```
|consistency_group_failover.status.message
|string
|query
|False
a|Filter by consistency_group_failover.status.message

* Introduced in: 9.8


|consistency_group_failover.status.code
|string
|query
|False
a|Filter by consistency_group_failover.status.code

* Introduced in: 9.8


|consistency_group_failover.type
|string
|query
|False
a|Filter by consistency_group_failover.type

* Introduced in: 9.14


|group_type
|string
|query
|False
a|Filter by group_type

* Introduced in: 9.11


|master_bias_activated_site
|string
|query
|False
a|Filter by master_bias_activated_site

* Introduced in: 9.14
```

```
|destination.luns.uuid
|string
|query
|False
a|Filter by destination.luns.uuid

* Introduced in: 9.16


|destination.luns.name
|string
|query
|False
a|Filter by destination.luns.name

* Introduced in: 9.16


|destination.svm.name
|string
|query
|False
a|Filter by destination.svm.name


|destination.svm.uuid
|string
|query
|False
a|Filter by destination.svm.uuid


|destination.consistency_group_volumes.name
|string
|query
|False
a|Filter by destination.consistency_group_volumes.name

* Introduced in: 9.8


|destination.cluster.name
|string
|query
|False
a|Filter by destination.cluster.name
```

```
* Introduced in: 9.7


|destination.cluster.uuid
|string
|query
|False
a|Filter by destination.cluster.uuid

* Introduced in: 9.7


|destination.path
|string
|query
|False
a|Filter by destination.path


|io_serving_copy
|string
|query
|False
a|Filter by io_serving_copy

* Introduced in: 9.14


|throttle
|integer
|query
|False
a|Filter by throttle

* Introduced in: 9.11


|last_transfer_type
|string
|query
|False
a|Filter by last_transfer_type

* Introduced in: 9.11


|preferred_site
```

```
|string
|query
|False
a|Filter by preferred_site

* Introduced in: 9.14


|exported_snapshot
|string
|query
|False
a|Filter by exported_snapshot


|identity_preservation
|string
|query
|False
a|Filter by identity_preservation

* Introduced in: 9.11


|unhealthy_reason.message
|string
|query
|False
a|Filter by unhealthy_reason.message


|unhealthy_reason.arguments
|string
|query
|False
a|Filter by unhealthy_reason.arguments

* Introduced in: 9.14


|unhealthy_reason.code
|string
|query
|False
a|Filter by unhealthy_reason.code
```

```
|svmdr_volumes.name
|string
|query
|False
a|Filter by svmdr_volumes.name

* Introduced in: 9.13


|total_transfer_duration
|string
|query
|False
a|Filter by total_transfer_duration

* Introduced in: 9.13


|healthy
|boolean
|query
|False
a|Filter by healthy


|uuid
|string
|query
|False
a|Filter by uuid


|total_transfer_bytes
|integer
|query
|False
a|Filter by total_transfer_bytes

* Introduced in: 9.13


|backoff_level
|string
|query
|False
a|Filter by backoff_level
```

```
* Introduced in: 9.14


|transfer_schedule.name
|string
|query
|False
a|Filter by transfer_schedule.name

* Introduced in: 9.11


|transfer_schedule.uuid
|string
|query
|False
a|Filter by transfer_schedule.uuid

* Introduced in: 9.11


|last_transfer_network_compression_ratio
|string
|query
|False
a|Filter by last_transfer_network_compression_ratio

* Introduced in: 9.13


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
```

```
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
```

```
a|Number of records


|records
|array[link:#snapmirror_relationship[snapmirror_relationship]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "backoff_level": "medium",
      "consistency_group_failover": {
        "error": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "state": "string",
        "status": {
          "code": "string",
```

```
          "message": "string"
        },
        "type": "string"
      },
      "destination": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "consistency_group_volumes": [
          {
            "name": "volume1"
          }
        ],
        "luns": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "path": "svm1:volume1",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "exported_snapshot": "string",
      "group_type": "consistency_group",
      "identity_preservation": "string",
      "io_serving_copy": "C1_sti85-vsim-ucs209a_cluster, C1_sti85-vsim-
ucs209c_cluster",
      "lag_time": "PT8H35M42S",
      "last_transfer_network_compression_ratio": 61,
```

```
      "last_transfer_type": "initialize",
      "master_bias_activated_site": "C1_sti85-vsim-ucs209a_cluster",
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "Asynchronous",
        "type": "string",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "preferred_site": "C1_sti85-vsim-ucs209a_cluster",
      "restore_to_snapshot": "string",
      "source": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "consistency_group_volumes": [
          {
            "name": "volume1"
          }
        ],
        "luns": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "/vol/volume1/lun1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "path": "svm1:volume1",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
```

```
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "state": "snapmirrored",
      "svmdr_volumes": [
        {
          "name": "volume1"
        }
      ],
      "throttle": 0,
      "total_transfer_bytes": 1098210312,
      "total_transfer_duration": "PT3M21S",
      "transfer": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "end_time": "2020-12-02 18:36:19 -0800",
        "last_updated_time": "2023-09-14 15:39:19 -0700",
        "state": "string",
        "total_duration": "PT28M41S",
        "type": "initialize",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "transfer_schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "unhealthy_reason": [
        {
          "arguments": [],
          "code": "6621444",
          "message": "Failed to complete update operation on one or more
item relationships."
        },
        {
          "arguments": [],
          "code": "6621445",
          "message": "Group Update failed"
        }
```

```
      ],
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303825
| Could not retrieve information for the SnapMirror policy type

| 13303817
| Unknown value for the Snapmirror State
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
```

```
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#error]
[.api-collapsible-fifth-title]
error

SnapMirror Consistency Group failover error message.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#status]
[.api-collapsible-fifth-title]
status

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Status code
```

```
|message
|string
a|SnapMirror Consistency Group failover status.


|===


[#snapmirror_consistency_group_failover]
[.api-collapsible-fifth-title]
snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|SnapMirror Consistency Group failover error message.


|state
|string
a|SnapMirror Consistency Group failover state.


|status
|link:#status[status]
a|

|type
|string
a|SnapMirror Consistency Group failover type.
```

```
|===


[#bucket_retention]
[.api-collapsible-fifth-title]
bucket_retention

Specifies the retention-mode and default retention period configured on
the destination bucket.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_period
|string
a|Specifies the default retention period that is applied to objects while
committing them to the WORM state without an associated retention period.
The retention period can be in years, or days. The retention period value
represents a duration and must be specified in the ISO-8601 duration
format.  A period specified for years and days is represented in the ISO-
8601 format as "P+++<num>+++Y" and "P+++<num>+++D" respectively, for
example "P10Y" represents a duration of 10 years. The period string must
contain only a single time element that is, either years, or days. A
duration which combines different periods is not supported, for example
"P1Y10D" is not supported.+++</num>++++++</num>+++


|mode
|string
a|The lock mode of the bucket.
compliance &dash; A SnapLock Compliance (SLC) bucket provides the highest
level of WORM protection and an administrator cannot destroy a compliance
bucket if it contains unexpired WORM objects.
 governance &dash; An administrator can delete a Governance bucket.
 no_lock &dash; Indicates the bucket does not support object locking.


|===


[#storage_service]
[.api-collapsible-fifth-title]
```

```
storage_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|This property indicates whether to create the destination endpoint using
storage service. This property is supported for Unified ONTAP destination
endpoints only.

* Introduced in: 9.7
* x-nullable: true


|enforce_performance
|boolean
a|Optional property to enforce storage service performance on the
destination endpoint. This property is applicable to FlexVol volume,
FlexGroup volume, and Consistency Group endpoints. This property is
supported for Unified ONTAP destination endpoints only.

* Default value:
* Introduced in: 9.7
* x-nullable: true


|name
|string
a|Optional property to specify the storage service name for the
destination endpoint. This property is considered when the property
"create_destination.storage_service.enabled" is set to "true". When the
property "create_destination.storage_service.enabled" is set to "true" and
the "create_destination.storage_service.name" for the endpoint is not
specified, then ONTAP selects the highest storage service available on the
cluster to provision the destination endpoint. This property is applicable
to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This
property is supported for Unified ONTAP destination endpoints only.

* enum: ["extreme", "performance", "value"]
* Introduced in: 9.6
* x-nullable: true
```

```
|===


[#tiering]
[.api-collapsible-fifth-title]
tiering

[cols=3*,options=header]
|===
|Name
|Type
|Description

|policy
|string
a|Optional property to specify the destination endpoint's tiering policy
when "create_destination.tiering.supported" is set to "true". This
property is applicable to FlexVol volume, FlexGroup volume, and
Consistency Group endpoints. This property determines whether the user
data blocks of the destination endpoint in a FabricPool will be tiered to
the cloud store when they become cold. FabricPool combines flash
(performance tier) with a cloud store into a single aggregate. Temperature
of the destination endpoint volume blocks increases if they are accessed
frequently and decreases when they are not.
all &dash; This policy allows tiering of both destination endpoint
snapshots and the user transferred data blocks to the cloud store as soon
as possible by ignoring the temperature on the volume blocks. This tiering
policy is not applicable for Consistency Group destination endpoints or
for synchronous relationships.
auto &dash; This policy allows tiering of both destination endpoint
snapshots and the active file system user data to the cloud store
none &dash; Destination endpoint volume blocks will not be tiered to the
cloud store.
snapshot_only &dash; This policy allows tiering of only the destination
endpoint volume snapshots not associated with the active file system. The
default tiering policy is "snapshot_only" for a FlexVol volume and "none"
for a FlexGroup volume. This property is supported for Unified ONTAP
destination endpoints only.

* enum: ["all", "auto", "none", "snapshot_only"]
* Introduced in: 9.6
* x-nullable: true


|supported
|boolean
```

```
a|Optional property to enable provisioning of the destination endpoint
volumes on FabricPool aggregates. This property is applicable to FlexVol
volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool
aggregates are used if this property is set to "true" and only non
FabricPool aggregates are used if this property is set to "false". Tiering
support for a FlexGroup volume can be changed by moving all of the
constituents to the required aggregates. Note that in order to tier data,
not only do the destination endpoint volumes need to support tiering by
using FabricPools, the "create_destination.tiering.policy" must not be
"none". A destination endpoint that uses FabricPools but has a tiering
"policy" of "none" supports tiering but will not tier any data. This
property is supported for Unified ONTAP destination endpoints only.

* Introduced in: 9.6
* x-nullable: true


|===


[#snapmirror_destination_creation]
[.api-collapsible-fifth-title]
snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a
SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-svm-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-svm-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-svm-name:> along with
the "destination.consistency_group_volumes" or "destination.luns" property
to indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
```

```
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.</destination-svm-name:></destination-svm-name:>+++</destination-svm-
name:dp-volume-name>+++


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bucket_retention
|link:#bucket_retention[bucket_retention]
a|Specifies the retention-mode and default retention period configured on
the destination bucket.


|enabled
|boolean
a|Optional property to create the destination endpoint when establishing a
SnapMirror relationship. It is assumed to be "false" if no other property
is set and assumed to be "true" if any other property is set.


|size
|integer
a|Optional property to specify the size of destination endpoint in bytes.
This property is applicable only to ONTAP S3 Bucket endpoints. The minimum
size for S3 bucket is 80MB and maximum size is 64TB. If not specified,
system will create destination with default size of 800GB.

* Introduced in: 9.10
* x-nullable: true


|snapshot_locking_enabled
|boolean
a|Optional property to create the destination endpoint with snapshot
locking enabled when establishing a SnapMirror relationship. This property
is applicable to FlexVol volumes and FlexGroup volumes.
```

```
|storage_service
|link:#storage_service[storage_service]
a|

|tiering
|link:#tiering[tiering]
a|

|===


[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#policy]
[.api-collapsible-fifth-title]
policy
```

```
Basic policy information of the relationship.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the SnapMirror policy.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#snapmirror_source_endpoint]
[.api-collapsible-fifth-title]
snapmirror_source_endpoint

Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.14
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#svmdr_volumes]
[.api-collapsible-fifth-title]
svmdr_volumes

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#transfer]
[.api-collapsible-fifth-title]
transfer

Basic information on the current transfer or the last transfer if there is
no active transfer at the time of the request.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|bytes_transferred
|integer
a|Total bytes transferred in the current or last successful transfer.


|end_time
|string
a|End time of the last transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in the current transfer.


|state
|string
a|
```

```
|total_duration
|string
a|Elapsed time to transfer all snapshots for the last successful transfer.


|type
|string
a|Specifies the operation type of the current transfer on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from "uninitialized" to "snapmirrored" or "in_sync". The _update_
transfer occurs when snapshots are being transferred from the source
endpoint to the destination endpoint as part of a scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from "broken_off" to "snapmirrored" or "in_sync". The _restore_ transfer
occurs when a snapshot is being restored from a destination endpoint to
another endpoint.


|uuid
|string
a|Transfer UUID. This property is applicable only for active transfers.


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_error]
[.api-collapsible-fifth-title]
snapmirror_error

SnapMirror error


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the error message encountered.


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#snapmirror_relationship]
[.api-collapsible-fifth-title]
```

```
snapmirror_relationship

SnapMirror relationship information. The SnapMirror relationship can be
either "async" or "sync" based on the type of SnapMirror policy associated
with the relationship. The source and destination endpoints of a
SnapMirror relationship must be of the same type, for example, if the
source endpoint is a FlexVol volume then the destination endpoint must be
a FlexVol volume.
The SnapMirror policy type "async" can be used when the SnapMirror
relationship has FlexVol volume or FlexGroup volume or SVM as the
endpoint. The SnapMirror policy type "sync" can be used when the
SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror
policy type "sync" with "sync_type" as "automated_failover" can be used
when the SnapMirror relationship has Consistency Group as the endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|consistency_group_failover
|link:#snapmirror_consistency_group_failover[snapmirror_consistency_group_
failover]
a|SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.

* Introduced in: 9.8
* readOnly: 1
```

```
|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
a|Snapshot exported to clients on destination.


|group_type
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
```

```
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
when the snapshot is restored from a destination endpoint to another
endpoint.


|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
```

```
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
PATCH state is being changed to "broken_off".


|restore
|boolean
a|Set to true to create a relationship for restore. To trigger restore-
transfer, use transfers POST on the restore relationship. SnapMirror
relationships with the policy type "async" can be restored. SnapMirror
relationships with the policy type "sync" cannot be restored.


|restore_to_snapshot
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.
```

```
|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
"sync". A PATCH call on the state change only triggers the transition to
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|svmdr_volumes
|array[link:#svmdr_volumes[svmdr_volumes]]
a|Specifies the list of constituent FlexVol volumes and FlexGroup volumes
for an SVM DR SnapMirror relationship. FlexGroup constituents are not
considered.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.
```

```
|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID55adfcabca56ada2f6fe9518bb67e043]]
= Create a SnapMirror relationship

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/snapmirror/relationships`#

*Introduced In:* 9.6

Creates a SnapMirror relationship. This API can optionally provision the
destination endpoint when it does not exist. This API must be executed on
the cluster containing the destination endpoint unless the destination
endpoint is being provisioned. When the destination endpoint is being
provisioned, this API can also be executed from the cluster containing the
```

```
source endpoint. Provisioning of the destination endpoint from the source
cluster is supported for the FlexVol volume, FlexGroup volume and
Application Consistency Group endpoints.

For SVM endpoints, provisioning the destination SVM endpoint is not
supported from the source cluster. When the destination endpoint exists,
the source SVM and the destination SVM must be in an SVM peer
relationship. When provisioning the destination endpoint, the SVM peer
relationship between the source SVM and the destination SVM is established
as part of the destination provision, provided that the source SVM has SVM
peering permissions for the destination cluster.

== Required properties

* `source.path` - Path to the source endpoint of the SnapMirror
relationship.
* `destination.path` - Path to the destination endpoint of the SnapMirror
relationship.
* `source.consistency_group_volumes` - List of FlexVol volumes of type
"RW" that are constituents of an Application Consistency Group.
* `destination.consistency_group_volumes` - List of FlexVol volumes of
type "DP" that are constituents of an Application Consistency Group.

== Recommended optional properties

* `policy.name` or `policy.uuid` - Policy governing the SnapMirror
relationship.
* `state` - Set the state to "snapmirrored" to automatically initialize
the relationship.
* `create_destination.enabled` - Enable this property to provision the
destination endpoint.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `policy.name` - _Asynchronous_
* `restore` - _false_
* `create_destination.tiering.policy` - `_snapshot_only_` (when
`create_destination.tiering.supported` is _true_ for FlexVol volume)
* `create_destination.tiering.policy` - `_none_` (when
`create_destination.tiering.supported` is _true_ for FlexGroup volume)
* `create_destination.storage_service.enforce_performance` - `_false_`
* `destination.ipspace` - `_Default_`
* `throttle` - _0_
```

```
* `backoff_level` - `_high_`
* `policy.name` - _Asynchronous_
* `restore` - _false_
* `destination.ipspace` - `_Default_`
* `throttle` - _0_
* `backoff_level` - `_high_`

== Related ONTAP commands

* `snapmirror create`
* `snapmirror protect`

== Important notes

* The property "transfer_schedule" if set on a SnapMirror relationship
overrides the "transfer_schedule" set on the policy being used with the
SnapMirror relationship.
* The property "throttle" if set on a SnapMirror relationship overrides
the "throttle" set on the policy being used with the SnapMirror
relationship.
* The properties "transfer_schedule" and "throttle" are not supported when
"restore" is set to "true".
* The property "transfer_schedule" cannot be set to null (no-quotes)
during SnapMirror relationship POST.
* The property "throttle" is not supported when
"create_destination.enabled" is set to "true".
* The property "identity_preservation" is applicable to only SnapMirror
relationships with SVM endpoints and it indicates which configuration of
the source SVM is replicated to the destination SVM.
* The property "backoff_level" is not supported when
"create_destination.enabled" is set to "true".
* The property "backoff_level" is only applicable for FlexVol SnapMirror
relationships.

== Examples

The following examples show how to create FlexVol volumes, FlexGroup
volumes, SVM and Application Consistency Group SnapMirror relationships.
Note that the source SVM name must be the local name of the peer SVM.

Creating a FlexVol SnapMirror relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"}}'
----
```

```
Creating a FlexGroup SnapMirror relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:source_flexgroup"}, "destination": { "path":
"dst_svm:dest_flexgroup"}}'
----

Creating a SVM SnapMirror relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": { "path": "src_svm:"},
"destination": { "path": "dst_svm:"}}'
----

Creating a SnapMirror relationship in order to restore from a destination.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"restore": "true"}'
----

Provision the destination FlexVol volume endpoint and create a SnapMirror
relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enabled": "true" }}'
----

Provision the destination FlexVol volume endpoint on a Fabricpool with a
tiering policy and create a SnapMirror relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enabled": "true", "tiering": { "supported":
"true", "policy": "auto" } } }'
----

Provision the destination FlexVol volume endpoint using storage service
and create a SnapMirror relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path":
```

```
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enabled": "true", "storage_service": { "enabled":
"true", "name": "extreme", "enforce_performance": "true" } } }'
----

Provision the destination SVM endpoint and create a SnapMirror
relationship of type XDP.

----
 POST "/api/snapmirror/relationships/" '{"source": {"path": "src_svm:",
"cluster": { "name": "cluster_src" }}, "destination": { "path":
"dst_svm:"}, "create_destination": { "enabled: "true" }}'
----

Create an asynchronous SnapMirror relationship with Application
Consistency Group endpoint.

----
 POST "/api/snapmirror/relationships/" '{"source": { "path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "policy": "Asynchronous" }'
----

Create an asynchronous SnapMirror relationship with a child Application
Consistency Group endpoint.

----
 POST "/api/snapmirror/relationships/" '{"source": { "path":
"src_svm:/cg/parent_src_cg/sub_child_cg" }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol" }, "policy": "MirrorAndVault" }'
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "create_destination": {
"enabled": "true", "tiering": { "supported": "true" } }, "policy":
"Asynchronous", "state": "snapmirrored" }'

----
 Create a SnapMirror active sync relationship with the Application
Consistency Group endpoint.
----
```

```
POST "/api/snapmirror/relationships/" '{"source": { "path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "policy": "AutomatedFailOver"
}'

----
 Provision the destination Application Consistency Group endpoint on a
Fabricpool with a tiering policy, create a SnapMirror active sync
relationship with a SnapMirror policy of type "sync" and sync_type of
"automated_failover", and initialize the SnapMirror relationship with
state as "in_sync".
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "create_destination": {
"enabled": "true", "tiering": { "supported": "true" } }, "policy":
"AutomatedFailOver", "state": "in_sync" }'

----
 Provision the destination Application Consistency Group endpoint with
storage service, create a SnapMirror active sync relationship with a
SnapMirror policy of type "sync" and sync_type of "automated_failover",
and initialize the SnapMirror relationship with state as "in_sync".
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "create_destination": {
"enabled": "true", "storage_service": { "enabled": "true", "name":
"extreme", "enforce_performance": "true" } }, "policy":
"AutomatedFailOver", "state": "in_sync" }'

----



Provision the destination Application Consistency Group endpoint with
storage service, create an asynchronous application consistency group
relationship with a SnapMirror policy of type "async" and an async_type of
```

```
"XDPDefault", and initialize the SnapMirror relationship with state as
"SnapMirrored".
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": [{ "name":
"src_vol_1" }, { "name": "src_vol_2" }] }, "destination": { "path":
"dst_svm:/cg/cg_dst_vol", "consistency_group_volumes": [{ "name":
"dst_vol_1" }, { "name": "dst_vol_2" }] }, "create_destination": {
"enabled": "true", "storage_service": { "enabled": "true", "name":
"extreme", "enforce_performance": "true" } }, "policy": "XDPDefault",
"state": "snapmirrored" }'

----



 Creating a FlexVol volume SnapMirror relationship of type XDP with
transfer_schedule and throttle.
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"transfer_schedule":{"uuid":"817500fa-092d-44c5-9c10-7b54f7b2f20a",
"name":"5min"}, "throttle":100}'

----




 Creating an asynchronous SnapMirror relationship with backoff_level set
to medium.
----

POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"backoff_level": "medium"}'
 ```

== Learn more

* xref:{relative_path}snapmirror_relationships_endpoint_overview.html[DOC
/snapmirror/relationships]


== Parameters
```
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|validate_only
|boolean
|query
|False
a|Validate the operation and its parameters, without actually performing
the operation.

* Introduced in: 9.7

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|create_destination
|link:#snapmirror_destination_creation[snapmirror_destination_creation]
a|Use this object to provision the destination endpoint when establishing
a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-SVM-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-SVM-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-SVM-
name:/cg/consistency-group-name> along with the
"destination.consistency_group_volumes" or "destination.luns" property to
indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
```

```
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.+++</destination-SVM-name:dp-volume-name>+++

* Introduced in: 9.7
* x-ntap-createOnly: true


|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
a|Snapshot exported to clients on destination.


|group_type
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
```

```
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
when the snapshot is restored from a destination endpoint to another
endpoint.


|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
```

```
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
PATCH state is being changed to "broken_off".


|restore
|boolean
a|Set to true to create a relationship for restore. To trigger restore-
transfer, use transfers POST on the restore relationship. SnapMirror
relationships with the policy type "async" can be restored. SnapMirror
relationships with the policy type "sync" cannot be restored.


|restore_to_snapshot
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
```

```
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
"sync". A PATCH call on the state change only triggers the transition to
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|svmdr_volumes
|array[link:#svmdr_volumes[svmdr_volumes]]
a|Specifies the list of constituent FlexVol volumes and FlexGroup volumes
for an SVM DR SnapMirror relationship. FlexGroup constituents are not
considered.
```

```
|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.


|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "backoff_level": "medium",
  "create_destination": {
    "bucket_retention": {
      "default_period": "P10Y",
      "mode": "governance"
    },
    "storage_service": {
      "name": "string"
    },
    "tiering": {
      "policy": "string"
    }
  },
  "destination": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ],
    "ipspace": "Default",
    "luns": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "group_type": "consistency_group",
  "identity_preservation": "string",
  "io_serving_copy": "C1_sti85-vsim-ucs209a_cluster, C1_sti85-vsim-
```

```
ucs209c_cluster",
  "lag_time": "PT8H35M42S",
  "last_transfer_network_compression_ratio": 61,
  "last_transfer_type": "initialize",
  "master_bias_activated_site": "C1_sti85-vsim-ucs209a_cluster",
  "policy": {
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "preferred_site": "C1_sti85-vsim-ucs209a_cluster",
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ],
    "luns": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "snapmirrored",
  "svmdr_volumes": [
    {
      "name": "volume1"
    }
  ],
  "throttle": 0,
  "total_transfer_bytes": 1098210312,
  "total_transfer_duration": "PT3M21S",
  "transfer": {
    "end_time": "2020-12-02 18:36:19 -0800",
    "last_updated_time": "2023-09-14 15:39:19 -0700",
    "state": "string",
```

```
    "total_duration": "PT28M41S",
    "type": "initialize",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unhealthy_reason": [
    {
      "arguments": [],
      "code": "6621444",
      "message": "Failed to complete update operation on one or more item
relationships."
    },
    {
      "arguments": [],
      "code": "6621445",
      "message": "Group Update failed"
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```


Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|918657
//end row
//start row
|Invalid duration value for the "retention_period" field.
//end row
//start row
|
//end row
//start row
|1115542
//end row
//start row
|Invalid value given for the field.
//end row
//start row
|
//end row
//start row
|6619441
//end row
//start row
|The source volume cannot be the same as the destination volume.
//end row
//start row
|
//end row
//start row
|6619599
//end row
//start row
|Only the "none" tiering policy is supported when creating a destination
volume with "snapshot_locking_enabled" set to true or if
```

```
"snapshot_locking_enabled" is set to true on the source volume.
//end row
//start row
|
//end row
//start row
|6619637
//end row
//start row
|Relationship with specified destination volume already exists.
//end row
//start row
|
//end row
//start row
|6619699
//end row
//start row
|Schedule not found.
//end row
//start row
|
//end row
//start row
|6620374
//end row
//start row
|Internal error. Failed to get SVM information.
//end row
//start row
|
//end row
//start row
|6620478
//end row
//start row
|Internal error. Failed to check SnapMirror capability.
//end row
//start row
|
//end row
//start row
|6620608
//end row
//start row
|Child consistency group cannot be a destination endpoint for SnapMirror
```

```
relationship creation.
//end row
//start row
|
//end row
//start row
|6620654
//end row
//start row
|Invalid SnapMirror Consistency Group name.
//end row
//start row
|
//end row
//start row
|6621125
//end row
//start row
|The policy is not valid for relationships with FlexGroup volume
endpoints. Only policies without snapshot creation schedules are supported
for these relationships.
//end row
//start row
|
//end row
//start row
|6621128
//end row
//start row
|Policy is not valid because it has snapshot creation schedules associated
with it. Such policies cannot be used for relationships with Consistency
Group endpoints.
//end row
//start row
|
//end row
//start row
|6621458
//end row
//start row
|The destination Consistency Group is the source of a SnapMirror
Synchronous (SM-S) relationship. Sources of SM-S relationships cannot be
the destination of any other SnapMirror relationship.
//end row
//start row
|
```

```
//end row
//start row
|6621782
//end row
//start row
|A property of the policy is not valid for relationships between these
endpoints.
//end row
//start row
|
//end row
//start row
|6621834
//end row
//start row
|Object store configuration does not exist for the specified SVM.
//end row
//start row
|
//end row
//start row
|6622088
//end row
//start row
|SnapMirror Asynchronous relationship is not supported on a Consistency
Group volume that has snapshot locking enabled.
//end row
//start row
|
//end row
//start row
|13303819
//end row
//start row
|Could not retrieve SnapMirror policy information.
//end row
//start row
|
//end row
//start row
|13303821
//end row
//start row
|Invalid SnapMirror policy UUID.
//end row
//start row
```

```
|
//end row
//start row
|13303841
//end row
//start row
|This operation is not supported for SnapMirror relationships between
these endpoints.
//end row
//start row
|
//end row
//start row
|13303852
//end row
//start row
|destination.path provided does not contain ":".
//end row
//start row
|
//end row
//start row
|13303853
//end row
//start row
|Restore relationships are not supported for SVM-DR endpoints.
//end row
//start row
|
//end row
//start row
|13303866
//end row
//start row
|Associating the specified SnapMirror policy with this SnapMirror
relationship is not supported.
//end row
//start row
|
//end row
//start row
|13303868
//end row
//start row
|Create of destination endpoint and SnapMirror relationship failed.
//end row
```

```
//start row
|
//end row
//start row
|13303869
//end row
//start row
|Creating a destination endpoint is not supported for restore
relationships.
//end row
//start row
|
//end row
//start row
|13303870
//end row
//start row
|A tiering policy cannot be specified if tiering is not being set to
supported.
//end row
//start row
|
//end row
//start row
|13303871
//end row
//start row
|Storage service properties cannot be specified if the storage service is
not being enabled.
//end row
//start row
|
//end row
//start row
|13303872
//end row
//start row
|Specified property requires a later effective cluster version.
//end row
//start row
|
//end row
//start row
|13303873
//end row
//start row
```

```
|Specifying a state when creating a relationship is only supported when
creating a destination endpoint.
//end row
//start row
|
//end row
//start row
|13303874
//end row
//start row
|Specified state is not supported when creating this relationship.
//end row
//start row
|
//end row
//start row
|13303875
//end row
//start row
|Destination aggregates do not have sufficient space for hosting copies of
source volumes.
//end row
//start row
|
//end row
//start row
|13303876
//end row
//start row
|Destination cluster does not have composite aggregates.
//end row
//start row
|
//end row
//start row
|13303877
//end row
//start row
|Source or destination cluster must be specified.
//end row
//start row
|
//end row
//start row
|13303878
//end row
```

```
//start row
|The specified fields do not match.
//end row
//start row
|
//end row
//start row
|13303879
//end row
//start row
|Source cluster name or UUID is needed to provision a destination SVM on
the local cluster.
//end row
//start row
|
//end row
//start row
|13303880
//end row
//start row
|Source cluster must be remote for provisioning a destination SVM on the
local cluster.
//end row
//start row
|
//end row
//start row
|13303881
//end row
//start row
|Network validation failed.
//end row
//start row
|
//end row
//start row
|13303882
//end row
//start row
|SVM validation failed.
//end row
//start row
|
//end row
//start row
|13303883
```

```
//end row
//start row
|Encryption is not enabled on the destination cluster.
//end row
//start row
|
//end row
//start row
|13303886
//end row
//start row
|SVM peer permission not found.
//end row
//start row
|
//end row
//start row
|13303887
//end row
//start row
|Synchronous SnapMirror relationships between FlexGroup volumes are not
supported.
//end row
//start row
|
//end row
//start row
|13303888
//end row
//start row
|Synchronous SnapMirror relationships require an effective cluster version
of 9.5 or later on both the source and destination clusters.
//end row
//start row
|
//end row
//start row
|13303889
//end row
//start row
|Asynchronous SnapMirror relationships between FlexGroup volumes require
an effective cluster version of 9.5 or later on both the source and
destination clusters.
//end row
//start row
|
```

```
//end row
//start row
|13303890
//end row
//start row
|Asynchronous SnapMirror relationships between FlexVol volumes require an
effective cluster version of 9.3, 9.5, or later on both the source and
destination clusters.
//end row
//start row
|
//end row
//start row
|13303891
//end row
//start row
|Creating a destination endpoint with storage service requires an
effective cluster version of 9.7 or later.
//end row
//start row
|
//end row
//start row
|13303892
//end row
//start row
|Fetching remote information from the destination cluster failed.
//end row
//start row
|
//end row
//start row
|13303893
//end row
//start row
|Updating job description failed.
//end row
//start row
|
//end row
//start row
|13303894
//end row
//start row
|Destination volume name is invalid. It must contain the source volume
name and have a suffix when creating a destination endpoint on a cluster
```

```
with an effective cluster version of 9.6 or earlier.
//end row
//start row
|
//end row
//start row
|13303895
//end row
//start row
|Operation on the remote destination cluster is not supported.
//end row
//start row
|
//end row
//start row
|13303916
//end row
//start row
|FlexGroup volumes are not supported on SnapLock aggregates.
//end row
//start row
|
//end row
//start row
|13303918
//end row
//start row
|No suitable destination aggregate type is available.
//end row
//start row
|
//end row
//start row
|13303919
//end row
//start row
|Only FabricPool enabled aggregates are available on the destination.
//end row
//start row
|
//end row
//start row
|13303920
//end row
//start row
|Only SnapLock aggregates are available on the destination. FlexGroup
```

```
volumes are not supported on SnapLock aggregates.
//end row
//start row
|
//end row
//start row
|13303921
//end row
//start row
|Unable to retrieve the SnapMirror capabilities of the destination
cluster.
//end row
//start row
|
//end row
//start row
|13303922
//end row
//start row
|Specified source SVM is not a data SVM.
//end row
//start row
|
//end row
//start row
|13303923
//end row
//start row
|Specified destination SVM is not a data SVM.
//end row
//start row
|
//end row
//start row
|13303924
//end row
//start row
|Source SVM has an invalid snapshot policy.
//end row
//start row
|
//end row
//start row
|13303925
//end row
//start row
```

```
|SnapMirror validation has failed.
//end row
//start row
|
//end row
//start row
|13303930
//end row
//start row
|The specified tiering policy is not supported for destination volumes of
Synchronous relationships.
//end row
//start row
|
//end row
//start row
|13303938
//end row
//start row
|Fetching information from the local cluster failed.
//end row
//start row
|
//end row
//start row
|13303939
//end row
//start row
|Could not create an SVM peer relationship.
//end row
//start row
|
//end row
//start row
|13303944
//end row
//start row
|An SVM-DR relationship is not supported because the source SVM has CIFS
configured and the associated SnapMirror policy has either the
"identity_preservation" property not set or set to
"exclude_network_and_protocol_config".
//end row
//start row
|
//end row
//start row
```

```
|13303949
//end row
//start row
|This SnapMirror policy is only supported for relationships with object
store destination endpoints.
//end row
//start row
|
//end row
//start row
|13303966
//end row
//start row
|Consistency Group relationships require a policy of type "sync" with a
sync_type of "automated_failover".
//end row
//start row
|
//end row
//start row
|13303967
//end row
//start row
|Consistency Group volume is not a FlexVol volume.
//end row
//start row
|
//end row
//start row
|13303968
//end row
//start row
|Unsupported volume type for the Consistency Group.
//end row
//start row
|
//end row
//start row
|13303969
//end row
//start row
|SnapMirror relationships between SVM endpoints and object store endpoints
are not supported.
//end row
//start row
|
```

```
//end row
//start row
|13303970
//end row
//start row
|Unsupported policy type for the Consistency Group.
//end row
//start row
|
//end row
//start row
|13303971
//end row
//start row
|SnapMirror relationships between Consistency Group endpoints and object
store endpoints are not supported.
//end row
//start row
|
//end row
//start row
|13303976
//end row
//start row
|Source or destination SVM is already part of an SVM-DR relation.
//end row
//start row
|
//end row
//start row
|13303977
//end row
//start row
|Destination Consistency Group volume UUIDs are not expected while
provisioning the destination volumes.
//end row
//start row
|
//end row
//start row
|13303978
//end row
//start row
|Number of Consistency Group volume names and UUIDs does not match.
//end row
//start row
```

```
|
//end row
//start row
|13303979
//end row
//start row
|Number of Consistency Group volumes exceeds the allowed limit.
//end row
//start row
|
//end row
//start row
|13303980
//end row
//start row
|Number of source and destination Consistency Group volumes do not match.
//end row
//start row
|
//end row
//start row
|13303981
//end row
//start row
|ISCSI or FCP protocol is not configured.
//end row
//start row
|
//end row
//start row
|13303982
//end row
//start row
|SAN data interface is not configured on the SVM.
//end row
//start row
|
//end row
//start row
|13304021
//end row
//start row
|No suitable storage can be found meeting the specified requirements. No
FabricPool enabled aggregates are available on the destination.
//end row
//start row
```

```
|
//end row
//start row
|13304022
//end row
//start row
|No suitable storage can be found meeting the specified requirements. No
non-root, non-taken-over, non-SnapLock, non-composite aggregates are
available on the destination.
//end row
//start row
|
//end row
//start row
|13304032
//end row
//start row
|In an "All SAN Array", an SVM-DR relationship is not supported when the
associated SnapMirror policy does not have the "identity_preservation"
property set to "exclude_network_and_protocol_config".
//end row
//start row
|
//end row
//start row
|13304080
//end row
//start row
|Specified UUID and name do not match.
//end row
//start row
|
//end row
//start row
|13304082
//end row
//start row
|Specified properties are mutually exclusive.
//end row
//start row
|
//end row
//start row
|13304083
//end row
//start row
```

```
|The specified property is not supported because all nodes in the cluster
are not capable of supporting the property.
//end row
//start row
|
//end row
//start row
|13304093
//end row
//start row
|The property specified is not supported for the specified relationships.
//end row
//start row
|
//end row
//start row
|13304098
//end row
//start row
|This SnapMirror policy is not supported for SnapMirror relationships with
SnapLock volumes.
//end row
//start row
|
//end row
//start row
|13304099
//end row
//start row
|SnapLock Compliance Clock is not running on all nodes in the destination
cluster.
//end row
//start row
|
//end row
//start row
|13304108
//end row
//start row
|Schedule not found in the Administrative SVM or the SVM for the
relationship.
//end row
//start row
|
//end row
//start row
```

```
|13304112
//end row
//start row
|File restore from a Consistency Group asynchronous SnapMirror
relationship endpoint is not supported.
//end row
//start row
|
//end row
//start row
|13304132
//end row
//start row
|Creating a destination endpoint is not supported with the "backoff_level"
property.
//end row
//start row
|
//end row
//start row
|13304144
//end row
//start row
|SnapLock is not licensed on the destination cluster.
//end row
//start row
|
//end row
//start row
|13304155
//end row
//start row
|Invalid SnapMirror consistency group name. The name must be from 1 to 30
characters in length. The first letter must be one of A-Z, a-z or "_". The
remaining characters must be A-Z, a-z, 0-9 or "_".
//end row
//start row
|
//end row
//start row
|13304160
//end row
//start row
|A relationship with the specified object store endpoint exists and a
restore relationship cannot be created with the same object store source
endpoint name as the destination of a backup relationship except when
```

```
restoring to the source volume of the existing relationship. The
"source.uuid", not the source endpoint name, establishes what object is
restored from the cloud.</private>
//end row
//start row
|
//end row
//start row
|53411897
//end row
//start row
|The specified source volumes do not match the volumes contained in the
source consistency group.
//end row
//start row
|Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.

===== Above error msgs have been added in sorted(ascending) order

===== of error codes, for example - 13303945 is error code for

===== SM_REST_SCHEDULE_CONFLICT_FOR_SM_OP and for 13303949 is err

===== code for SM_REST_POLICY_ONLY_SUPPORTED_FOR_OBJSTORE, so

===== SM_REST_SCHEDULE_CONFLICT_FOR_SM_OP has been placed before

===== SM_REST_POLICY_ONLY_SUPPORTED_FOR_OBJSTORE.

===== Also for private error msg, add private tags as  after adding the
error msg at correct place(in sorted order).

===== Please make sure new error messages are being put at correct place

===== so that order is maintained.

//end row
|===
//end table


== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
```

```
[.api-collapsible-fifth-title]
error

SnapMirror Consistency Group failover error message.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#status]
[.api-collapsible-fifth-title]
status

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Status code


|message
|string
```

```
a|SnapMirror Consistency Group failover status.


|===


[#snapmirror_consistency_group_failover]
[.api-collapsible-fifth-title]
snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|SnapMirror Consistency Group failover state.


|status
|link:#status[status]
a|

|type
|string
a|SnapMirror Consistency Group failover type.


|===


[#bucket_retention]
[.api-collapsible-fifth-title]
bucket_retention

Specifies the retention-mode and default retention period configured on
```

```
the destination bucket.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_period
|string
a|Specifies the default retention period that is applied to objects while
committing them to the WORM state without an associated retention period.
The retention period can be in years, or days. The retention period value
represents a duration and must be specified in the ISO-8601 duration
format.  A period specified for years and days is represented in the ISO-
8601 format as "P+++<num>+++Y" and "P+++<num>+++D" respectively, for
example "P10Y" represents a duration of 10 years. The period string must
contain only a single time element that is, either years, or days. A
duration which combines different periods is not supported, for example
"P1Y10D" is not supported.+++</num>++++++</num>+++


|mode
|string
a|The lock mode of the bucket.
compliance &dash; A SnapLock Compliance (SLC) bucket provides the highest
level of WORM protection and an administrator cannot destroy a compliance
bucket if it contains unexpired WORM objects.
 governance &dash; An administrator can delete a Governance bucket.
 no_lock &dash; Indicates the bucket does not support object locking.


|===


[#storage_service]
[.api-collapsible-fifth-title]
storage_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
```

```
|boolean
a|This property indicates whether to create the destination endpoint using
storage service. This property is supported for Unified ONTAP destination
endpoints only.

* Introduced in: 9.7
* x-nullable: true


|enforce_performance
|boolean
a|Optional property to enforce storage service performance on the
destination endpoint. This property is applicable to FlexVol volume,
FlexGroup volume, and Consistency Group endpoints. This property is
supported for Unified ONTAP destination endpoints only.

* Default value: 1
* Introduced in: 9.7
* x-nullable: true


|name
|string
a|Optional property to specify the storage service name for the
destination endpoint. This property is considered when the property
"create_destination.storage_service.enabled" is set to "true". When the
property "create_destination.storage_service.enabled" is set to "true" and
the "create_destination.storage_service.name" for the endpoint is not
specified, then ONTAP selects the highest storage service available on the
cluster to provision the destination endpoint. This property is applicable
to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This
property is supported for Unified ONTAP destination endpoints only.

* enum: ["extreme", "performance", "value"]
* Introduced in: 9.6
* x-nullable: true


|===


[#tiering]
[.api-collapsible-fifth-title]
tiering

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|policy
|string
a|Optional property to specify the destination endpoint's tiering policy
when "create_destination.tiering.supported" is set to "true". This
property is applicable to FlexVol volume, FlexGroup volume, and
Consistency Group endpoints. This property determines whether the user
data blocks of the destination endpoint in a FabricPool will be tiered to
the cloud store when they become cold. FabricPool combines flash
(performance tier) with a cloud store into a single aggregate. Temperature
of the destination endpoint volume blocks increases if they are accessed
frequently and decreases when they are not.
all &dash; This policy allows tiering of both destination endpoint
snapshots and the user transferred data blocks to the cloud store as soon
as possible by ignoring the temperature on the volume blocks. This tiering
policy is not applicable for Consistency Group destination endpoints or
for synchronous relationships.
auto &dash; This policy allows tiering of both destination endpoint
snapshots and the active file system user data to the cloud store
none &dash; Destination endpoint volume blocks will not be tiered to the
cloud store.
snapshot_only &dash; This policy allows tiering of only the destination
endpoint volume snapshots not associated with the active file system. The
default tiering policy is "snapshot_only" for a FlexVol volume and "none"
for a FlexGroup volume. This property is supported for Unified ONTAP
destination endpoints only.

* enum: ["all", "auto", "none", "snapshot_only"]
* Introduced in: 9.6
* x-nullable: true


|supported
|boolean
a|Optional property to enable provisioning of the destination endpoint
volumes on FabricPool aggregates. This property is applicable to FlexVol
volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool
aggregates are used if this property is set to "true" and only non
FabricPool aggregates are used if this property is set to "false". Tiering
support for a FlexGroup volume can be changed by moving all of the
constituents to the required aggregates. Note that in order to tier data,
not only do the destination endpoint volumes need to support tiering by
using FabricPools, the "create_destination.tiering.policy" must not be
```

```
"none". A destination endpoint that uses FabricPools but has a tiering
"policy" of "none" supports tiering but will not tier any data. This
property is supported for Unified ONTAP destination endpoints only.

* Introduced in: 9.6
* x-nullable: true


|===


[#snapmirror_destination_creation]
[.api-collapsible-fifth-title]
snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a
SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-svm-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-svm-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-svm-name:> along with
the "destination.consistency_group_volumes" or "destination.luns" property
to indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.</destination-svm-name:></destination-svm-name:>+++</destination-svm-
name:dp-volume-name>+++
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|bucket_retention
|link:#bucket_retention[bucket_retention]
a|Specifies the retention-mode and default retention period configured on
the destination bucket.


|enabled
|boolean
a|Optional property to create the destination endpoint when establishing a
SnapMirror relationship. It is assumed to be "false" if no other property
is set and assumed to be "true" if any other property is set.


|size
|integer
a|Optional property to specify the size of destination endpoint in bytes.
This property is applicable only to ONTAP S3 Bucket endpoints. The minimum
size for S3 bucket is 80MB and maximum size is 64TB. If not specified,
system will create destination with default size of 800GB.

* Introduced in: 9.10
* x-nullable: true


|snapshot_locking_enabled
|boolean
a|Optional property to create the destination endpoint with snapshot
locking enabled when establishing a SnapMirror relationship. This property
is applicable to FlexVol volumes and FlexGroup volumes.


|storage_service
|link:#storage_service[storage_service]
a|

|tiering
|link:#tiering[tiering]
a|

|===
```

```
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|uuid
|string
a|Unique identifier of the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* x-ntap-createOnly: true
* Introduced in: 9.8
* x-nullable: true
```

```
|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
```

```
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|ipspace
|string
a|Optional property to specify the IPSpace of the SVM.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#policy]
[.api-collapsible-fifth-title]
policy

Basic policy information of the relationship.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|Name of the SnapMirror policy.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#snapmirror_source_endpoint]
[.api-collapsible-fifth-title]
snapmirror_source_endpoint

Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
```

```
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.14
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#svmdr_volumes]
[.api-collapsible-fifth-title]
svmdr_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===
```

```
[#transfer]
[.api-collapsible-fifth-title]
transfer

Basic information on the current transfer or the last transfer if there is
no active transfer at the time of the request.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_transferred
|integer
a|Total bytes transferred in the current or last successful transfer.


|end_time
|string
a|End time of the last transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in the current transfer.


|state
|string
a|

|total_duration
|string
a|Elapsed time to transfer all snapshots for the last successful transfer.


|type
|string
a|Specifies the operation type of the current transfer on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from "uninitialized" to "snapmirrored" or "in_sync". The _update_
transfer occurs when snapshots are being transferred from the source
endpoint to the destination endpoint as part of a scheduled or manual
```

```
update. The _resync_ transfer occurs when the relationship state changes
from "broken_off" to "snapmirrored" or "in_sync". The _restore_ transfer
occurs when a snapshot is being restored from a destination endpoint to
another endpoint.


|uuid
|string
a|Transfer UUID. This property is applicable only for active transfers.


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_error]
[.api-collapsible-fifth-title]
snapmirror_error
```

```
SnapMirror error


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the error message encountered.


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#snapmirror_relationship]
[.api-collapsible-fifth-title]
snapmirror_relationship

SnapMirror relationship information. The SnapMirror relationship can be
either "async" or "sync" based on the type of SnapMirror policy associated
with the relationship. The source and destination endpoints of a
SnapMirror relationship must be of the same type, for example, if the
source endpoint is a FlexVol volume then the destination endpoint must be
a FlexVol volume.
The SnapMirror policy type "async" can be used when the SnapMirror
relationship has FlexVol volume or FlexGroup volume or SVM as the
endpoint. The SnapMirror policy type "sync" can be used when the
SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror
policy type "sync" with "sync_type" as "automated_failover" can be used
when the SnapMirror relationship has Consistency Group as the endpoint.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|create_destination
|link:#snapmirror_destination_creation[snapmirror_destination_creation]
a|Use this object to provision the destination endpoint when establishing
a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-SVM-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-SVM-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-SVM-
name:/cg/consistency-group-name> along with the
"destination.consistency_group_volumes" or "destination.luns" property to
indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.+++</destination-SVM-name:dp-volume-name>+++

* Introduced in: 9.7
* x-ntap-createOnly: true
```

```
|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
a|Snapshot exported to clients on destination.


|group_type
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
```

```
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
when the snapshot is restored from a destination endpoint to another
endpoint.


|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
```

```
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
PATCH state is being changed to "broken_off".


|restore
|boolean
a|Set to true to create a relationship for restore. To trigger restore-
transfer, use transfers POST on the restore relationship. SnapMirror
relationships with the policy type "async" can be restored. SnapMirror
relationships with the policy type "sync" cannot be restored.


|restore_to_snapshot
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
```

```
the remote endpoint.


|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
"sync". A PATCH call on the state change only triggers the transition to
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|svmdr_volumes
|array[link:#svmdr_volumes[svmdr_volumes]]
a|Specifies the list of constituent FlexVol volumes and FlexGroup volumes
for an SVM DR SnapMirror relationship. FlexGroup constituents are not
considered.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
```

```
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.


|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#job_link]
[.api-collapsible-fifth-title]
```

```
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID6de695e38491542c513c917b23be8f41]]
= Delete a SnapMirror relationship

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/snapmirror/relationships/{uuid}`#

*Introduced In:* 9.6

Deletes a SnapMirror relationship.

== Important notes

* The "destination_only", "source_only", and "source_info_only" flags are
mutually exclusive. If no flag is specified, the relationship is deleted
from both the source and destination and all common snapshots between the
source and destination are also deleted.
* For a restore relationship, the call must be executed on the cluster
containing the destination endpoint without specifying the
destination_only, source_only, or source_info_only parameters.
* Additionally, ensure that there are no ongoing transfers on a restore
relationship before calling this API.
* The "failover", "force-failover" and "failback" query parameters are
only applicable for SVM-DR SnapMirror relationships.
* When a SnapMirror relationship associated with a pair of source and
destination Consistency Groups is deleted, the corresponding Consistency
Groups on the source and destination clusters are not automatically
deleted and remain in place.
* The "delete_lun_maps_in_destination" query parameter is applicable only
for SnapMirror active sync relationships.
* The "unmap_namespace" query parameter is applicable only for NVMe
SnapMirror active sync relationships.

== Related ONTAP commands

* `snapmirror delete`
* `snapmirror release`

== Examples
```

```
The following examples show how to delete the relationship from both the
source and destination, the destination only, and the source only.

Deleting the relationship from both the source and destination. This API
must be run on the cluster containing the destination endpoint.

----
 DELETE "/api/snapmirror/relationships/4512b2d2-fd60-11e8-8929-
005056bbfe52"
----

Deleting the relationship on the destination only. This API must be run on
the cluster containing the destination endpoint.

----
 DELETE "/api/snapmirror/relationships/fd1e0697-02ba-11e9-acc7-
005056a7697f/?destination_only=true"
----

Deleting the relationship on the source only. This API must be run on the
cluster containing the source endpoint.

----
 DELETE "/api/snapmirror/relationships/93e828ba-02bc-11e9-acc7-
005056a7697f/?source_only=true"
----

Deleting the source information only. This API must be run on the cluster
containing the source endpoint. This does not delete the common snapshots
between the source and destination.

----
 DELETE "/api/snapmirror/relationships/caf545a2-fc60-11e8-aa13-
005056a707ff/?source_info_only=true"
----

Deleting the relationship from source and destination cluster along with
deleting the LUN maps for the volumes of the CG in destination cluster.
This API must be run on the cluster containing the destination endpoint.

----
 DELETE "/api/snapmirror/relationships/feda8f5e-e29e-11ed-94aa-
005056a78ce2/?delete_lun_maps_in_destination=true"
----

Deleting the relationship from destination cluster along with deleting the
namespace maps for the volumes of the CG in destination cluster. This API
```

```
must be run on the cluster containing the destination endpoint.

----
 DELETE "/api/snapmirror/relationships/feda8f5e-e29e-11ed-94aa-
005056a78ce2/?unmap_namespace=true"
----

== Learn more

* xref:{relative_path}snapmirror_relationships_endpoint_overview.html[DOC
/snapmirror/relationships]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|SnapMirror relationship UUID


|destination_only
|boolean
|query
|False
a|Deletes a relationship on the destination only. This parameter is
applicable only when the call is executed on the cluster that contains the
destination endpoint.


|source_only
|boolean
|query
|False
a|Deletes a relationship on the source only. This parameter is applicable
only when the call is executed on the cluster that contains the source
endpoint. In the case of a relationship that uses SnapMirror policy of
```

```
type 'continuous', this parameter is used to delete the source end of the
relationship while ignoring errors when the destination endpoint is not
reachable.


|source_info_only
|boolean
|query
|False
a|Deletes relationship information on the source only. This parameter is
applicable only when the call is executed on the cluster that contains the
source endpoint.


|delete_lun_maps_in_destination
|boolean
|query
|False
a|Deletes LUN mapping for the volumes of the CG in destination.

* Introduced in: 9.14
* Default value:


|unmap_namespace
|boolean
|query
|False
a|Deletes namespace mapping and then deletes the subsystems.

* Introduced in: 9.17
* Default value:


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
```

```
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303825
| Could not retrieve information for the SnapMirror policy type

| 13303814
| Could not retrieve the source or destination SVM UUID

| 13303815
| Could not retrieve information for the peer cluster

| 13303822
| SnapMirror release has failed

| 13303813
| SnapMirror release was successful but delete has failed

| 13303854
| Cleanup of restore relationship failed

| 13303855
| DELETE call on a restore relationship does not support the given flags

| 13303865
| Deleting the specified SnapMirror policy is not supported.

| 6619715
| Modification of relationship is in progress. Retry the command after a
few minutes.

| 6621470
| SnapMirror delete failed on the relationship because "unmap-namespace"
is not set. Set the "unmap-namespace" to true for NVMe SnapMirror active
sync relationships.
|===



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID99767c4b87ccc9cfbbb8af03f1476ee4]]
= Retrieve a SnapMirror relationship

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/relationships/{uuid}`#

*Introduced In:* 9.6

Retrieves a SnapMirror relationship.

== Related ONTAP commands

* `snapmirror show`
* `snapmirror list-destinations`

== Expensive properties

* `source.consistency_group_volumes.name`
* `destination.consistency_group_volumes.name`

== Example

----
GET "/api/snapmirror/relationships/caf545a2-fc60-11e8-aa13-005056a707ff/"
----

== Learn more

* xref:{relative_path}snapmirror_relationships_endpoint_overview.html[DOC
/snapmirror/relationships]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|SnapMirror relationship UUID


|list_destinations_only
|boolean
|query
|False
a|Set to true to show relationships from the source only.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|consistency_group_failover
|link:#snapmirror_consistency_group_failover[snapmirror_consistency_group_
failover]
a|SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.

* Introduced in: 9.8
* readOnly: 1


|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
a|Snapshot exported to clients on destination.


|group_type
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
```

```
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
when the snapshot is restored from a destination endpoint to another
endpoint.
```

```
|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
PATCH state is being changed to "broken_off".


|restore
|boolean
a|Set to true to create a relationship for restore. To trigger restore-
transfer, use transfers POST on the restore relationship. SnapMirror
```

```
relationships with the policy type "async" can be restored. SnapMirror
relationships with the policy type "sync" cannot be restored.


|restore_to_snapshot
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
```

```
"sync". A PATCH call on the state change only triggers the transition to
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|svmdr_volumes
|array[link:#svmdr_volumes[svmdr_volumes]]
a|Specifies the list of constituent FlexVol volumes and FlexGroup volumes
for an SVM DR SnapMirror relationship. FlexGroup constituents are not
considered.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.


|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
```

```
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "backoff_level": "medium",
  "consistency_group_failover": {
    "error": {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist"
    },
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
```

```
  "destination": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1"
      }
    ],
    "luns": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "group_type": "consistency_group",
  "identity_preservation": "string",
  "io_serving_copy": "C1_sti85-vsim-ucs209a_cluster, C1_sti85-vsim-
ucs209c_cluster",
  "lag_time": "PT8H35M42S",
  "last_transfer_network_compression_ratio": 61,
  "last_transfer_type": "initialize",
  "master_bias_activated_site": "C1_sti85-vsim-ucs209a_cluster",
  "policy": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "preferred_site": "C1_sti85-vsim-ucs209a_cluster",
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1"
      }
    ],
    "luns": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "snapmirrored",
```

```
  "svmdr_volumes": [
    {
      "name": "volume1"
    }
  ],
  "throttle": 0,
  "total_transfer_bytes": 1098210312,
  "total_transfer_duration": "PT3M21S",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "end_time": "2020-12-02 18:36:19 -0800",
    "last_updated_time": "2023-09-14 15:39:19 -0700",
    "state": "string",
    "total_duration": "PT28M41S",
    "type": "initialize",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "transfer_schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unhealthy_reason": [
    {
      "arguments": [],
      "code": "6621444",
      "message": "Failed to complete update operation on one or more item
relationships."
    },
    {
      "arguments": [],
      "code": "6621445",
      "message": "Group Update failed"
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====
```

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303825
| Could not retrieve information for the SnapMirror policy type

| 13303817
| Unknown value for the Snapmirror State
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
```

```
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

SnapMirror Consistency Group failover error message.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|===


[#status]
[.api-collapsible-fifth-title]
status

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Status code


|message
|string
a|SnapMirror Consistency Group failover status.


|===


[#snapmirror_consistency_group_failover]
[.api-collapsible-fifth-title]
snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#error[error]
a|SnapMirror Consistency Group failover error message.


|state
|string
a|SnapMirror Consistency Group failover state.


|status
|link:#status[status]
a|

|type
|string
a|SnapMirror Consistency Group failover type.


|===


[#bucket_retention]
[.api-collapsible-fifth-title]
bucket_retention

Specifies the retention-mode and default retention period configured on
the destination bucket.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_period
|string
a|Specifies the default retention period that is applied to objects while
committing them to the WORM state without an associated retention period.
The retention period can be in years, or days. The retention period value
represents a duration and must be specified in the ISO-8601 duration
format.  A period specified for years and days is represented in the ISO-
8601 format as "P+++<num>+++Y" and "P+++<num>+++D" respectively, for
example "P10Y" represents a duration of 10 years. The period string must
contain only a single time element that is, either years, or days. A
```

```
duration which combines different periods is not supported, for example
"P1Y10D" is not supported.+++</num>++++++</num>+++


|mode
|string
a|The lock mode of the bucket.
compliance &dash; A SnapLock Compliance (SLC) bucket provides the highest
level of WORM protection and an administrator cannot destroy a compliance
bucket if it contains unexpired WORM objects.
 governance &dash; An administrator can delete a Governance bucket.
 no_lock &dash; Indicates the bucket does not support object locking.


|===


[#storage_service]
[.api-collapsible-fifth-title]
storage_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|This property indicates whether to create the destination endpoint using
storage service. This property is supported for Unified ONTAP destination
endpoints only.

* Introduced in: 9.7
* x-nullable: true


|enforce_performance
|boolean
a|Optional property to enforce storage service performance on the
destination endpoint. This property is applicable to FlexVol volume,
FlexGroup volume, and Consistency Group endpoints. This property is
supported for Unified ONTAP destination endpoints only.

* Default value: 1
* Introduced in: 9.7
* x-nullable: true
```

```
|name
|string
a|Optional property to specify the storage service name for the
destination endpoint. This property is considered when the property
"create_destination.storage_service.enabled" is set to "true". When the
property "create_destination.storage_service.enabled" is set to "true" and
the "create_destination.storage_service.name" for the endpoint is not
specified, then ONTAP selects the highest storage service available on the
cluster to provision the destination endpoint. This property is applicable
to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This
property is supported for Unified ONTAP destination endpoints only.

* enum: ["extreme", "performance", "value"]
* Introduced in: 9.6
* x-nullable: true


|===


[#tiering]
[.api-collapsible-fifth-title]
tiering

[cols=3*,options=header]
|===
|Name
|Type
|Description

|policy
|string
a|Optional property to specify the destination endpoint's tiering policy
when "create_destination.tiering.supported" is set to "true". This
property is applicable to FlexVol volume, FlexGroup volume, and
Consistency Group endpoints. This property determines whether the user
data blocks of the destination endpoint in a FabricPool will be tiered to
the cloud store when they become cold. FabricPool combines flash
(performance tier) with a cloud store into a single aggregate. Temperature
of the destination endpoint volume blocks increases if they are accessed
frequently and decreases when they are not.
all &dash; This policy allows tiering of both destination endpoint
snapshots and the user transferred data blocks to the cloud store as soon
as possible by ignoring the temperature on the volume blocks. This tiering
```

```
policy is not applicable for Consistency Group destination endpoints or
for synchronous relationships.
auto &dash; This policy allows tiering of both destination endpoint
snapshots and the active file system user data to the cloud store
none &dash; Destination endpoint volume blocks will not be tiered to the
cloud store.
snapshot_only &dash; This policy allows tiering of only the destination
endpoint volume snapshots not associated with the active file system. The
default tiering policy is "snapshot_only" for a FlexVol volume and "none"
for a FlexGroup volume. This property is supported for Unified ONTAP
destination endpoints only.

* enum: ["all", "auto", "none", "snapshot_only"]
* Introduced in: 9.6
* x-nullable: true


|supported
|boolean
a|Optional property to enable provisioning of the destination endpoint
volumes on FabricPool aggregates. This property is applicable to FlexVol
volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool
aggregates are used if this property is set to "true" and only non
FabricPool aggregates are used if this property is set to "false". Tiering
support for a FlexGroup volume can be changed by moving all of the
constituents to the required aggregates. Note that in order to tier data,
not only do the destination endpoint volumes need to support tiering by
using FabricPools, the "create_destination.tiering.policy" must not be
"none". A destination endpoint that uses FabricPools but has a tiering
"policy" of "none" supports tiering but will not tier any data. This
property is supported for Unified ONTAP destination endpoints only.

* Introduced in: 9.6
* x-nullable: true


|===


[#snapmirror_destination_creation]
[.api-collapsible-fifth-title]
snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a
SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
```

```
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-svm-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-svm-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-svm-name:> along with
the "destination.consistency_group_volumes" or "destination.luns" property
to indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.</destination-svm-name:></destination-svm-name:>+++</destination-svm-
name:dp-volume-name>+++


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bucket_retention
|link:#bucket_retention[bucket_retention]
a|Specifies the retention-mode and default retention period configured on
the destination bucket.


|enabled
|boolean
a|Optional property to create the destination endpoint when establishing a
SnapMirror relationship. It is assumed to be "false" if no other property
is set and assumed to be "true" if any other property is set.
```

```
|size
|integer
a|Optional property to specify the size of destination endpoint in bytes.
This property is applicable only to ONTAP S3 Bucket endpoints. The minimum
size for S3 bucket is 80MB and maximum size is 64TB. If not specified,
system will create destination with default size of 800GB.

* Introduced in: 9.10
* x-nullable: true


|snapshot_locking_enabled
|boolean
a|Optional property to create the destination endpoint with snapshot
locking enabled when establishing a SnapMirror relationship. This property
is applicable to FlexVol volumes and FlexGroup volumes.


|storage_service
|link:#storage_service[storage_service]
a|

|tiering
|link:#tiering[tiering]
a|

|===


[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|
```

```
|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of a LUN.
```

```
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapmirror_endpoint]
```

```
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
```

```
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#policy]
[.api-collapsible-fifth-title]
policy

Basic policy information of the relationship.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the SnapMirror policy.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===
```

```
[#snapmirror_source_endpoint]
[.api-collapsible-fifth-title]
snapmirror_source_endpoint

Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
```

```
* Introduced in: 9.14
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#svmdr_volumes]
[.api-collapsible-fifth-title]
svmdr_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#transfer]
[.api-collapsible-fifth-title]
transfer

Basic information on the current transfer or the last transfer if there is
no active transfer at the time of the request.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|bytes_transferred
|integer
a|Total bytes transferred in the current or last successful transfer.


|end_time
|string
a|End time of the last transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in the current transfer.


|state
|string
a|

|total_duration
|string
a|Elapsed time to transfer all snapshots for the last successful transfer.


|type
|string
a|Specifies the operation type of the current transfer on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from "uninitialized" to "snapmirrored" or "in_sync". The _update_
transfer occurs when snapshots are being transferred from the source
endpoint to the destination endpoint as part of a scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from "broken_off" to "snapmirrored" or "in_sync". The _restore_ transfer
occurs when a snapshot is being restored from a destination endpoint to
another endpoint.


|uuid
|string
a|Transfer UUID. This property is applicable only for active transfers.


|===


[#transfer_schedule]
```

```
[.api-collapsible-fifth-title]
transfer_schedule

Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_error]
[.api-collapsible-fifth-title]
snapmirror_error

SnapMirror error


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[string]
a|Arguments present in the error message encountered.


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID2e7b162aff2ac555cfa610e6cd6429d4]]
= Update a SnapMirror relationship

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/snapmirror/relationships/{uuid}`#

*Introduced In:* 9.6

Updates a SnapMirror relationship. This API is used to initiate SnapMirror
operations such as "initialize", "resync", "break", "quiesce", and
"resume" by specifying the appropriate value for the "state" field. It is
also used to modify the SnapMirror policy associated with the specified
relationship. Additionally, a SnapMirror relationship can be failed over
to the destination endpoint or a failed over SnapMirror relationship can
be failed back to the original state or a SnapMirror relationship
direction can be reversed using this API. This API can also be used to
expand the SnapMirror active sync relationship with the specified source
and destination volumes.

To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or "in_sync" for relationships
with a policy of type "sync".

To break the relationship or to failover to the destination endpoint and
start serving data from the destination endpoint, PATCH the state to
"broken_off" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and sync_type as
"automated_failover" cannot be "broken_off".

To resync the broken relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or "in_sync" for relationships
with a policy of type "sync".

To failback the failed over relationship and start serving data from the
source endpoint, PATCH the state to "snapmirrored" for relationships with
a policy of type "async" or "in_sync" for relationships with a policy of
type "sync" and set the query flag "failback" as "true". SnapMirror
relationships with the policy type as "sync" and sync_type as
"automated_failover" can be in "broken_off" state due to a failed attempt
```

```
of automated SnapMirror failover operation.

To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and sync_type as
"automated_failover" cannot be "paused".

To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or
"in_sync" for relationships with a policy of type "sync".

To reverse the direction of the relationship, PATCH the "source.path" with
the destination endpoint and the "destination.path" with the source
endpoint and the relationship state to "snapmirrored" for relationships
with a policy of type "async" or "in_sync" for relationships with a policy
of type "sync". For relationships with a policy of type "async" and
relationship state as "snapmirrored", stop IO on the source endpoint and
perform a SnapMirror transfer POST operation before reversing the
direction of the relationship to prevent any loss of data.

The values "in_sync", "out_of_sync", and "synchronizing" are only
applicable to relationships with a policy of type "sync".

When "transfer_schedule" is specified along with "state" during PATCH,
first the schedule is modified on the relationship and then the respective
SnapMirror operation is initiated. The "transfer_schedule" specified is
used to update asynchronous relationships.

When "throttle" is specified along with "state" during PATCH, first the
throttle is modified on the relationship, which will be used by any
upcoming transfers and then the respective SnapMirror operation is
initiated. If the SnapMirror operation initiated a transfer then it will
also use the new throttle. If "throttle" needs to be applied for a
specific transfer use SnapMirror Transfer REST API.

For SnapMirror active sync relationships, when "consistency_group_volumes"
is specified during PATCH, first the existing FlexVol volume relationship
is deleted and released and then the volumes are added to the SnapMirror
active sync relationship.

== Examples

== Related ONTAP commands

* `snapmirror modify`
* `snapmirror initialize`
* `snapmirror resync`
```

```
* `snapmirror break`
* `snapmirror quiesce`
* `snapmirror resume`

== Important notes

* The property "transfer_schedule" if set on a SnapMirror relationship
overrides the "transfer_schedule" set on the policy being used with the
SnapMirror relationship.
* The property "throttle" if set on a SnapMirror relationship overrides
the "throttle" set on the policy being used with the SnapMirror
relationship.
* The properties "transfer_schedule" and "throttle" are not supported when
"failback" is set to "true".
* The properties "transfer_schedule" and "throttle" are not supported when
"failover" is set to "true".
* The properties "transfer_schedule" and "throttle" are not supported when
"force_failover" is set to "true".
* The properties "transfer_schedule" and "throttle" are not supported when
the direction of the relationship is being reversed.
* To remove a transfer_schedule on a SnapMirror relationship set the
"transfer_schedule" to null (no-quotes) during SnapMirror relationship
PATCH.
* The property "identity_preservation" value can be changed from a higher
"identity_preservation" threshold value to a lower "identity_preservation"
threshold value but not vice-versa. For example, the threshold value of
the "identity_preservation" property can be changed from "full" to
"exclude_network_config", but cannot be increased from
"exclude_network_and_protocol_config" to "exclude_network_config" to
"full". The threshold value of the "identity_preservation" cannot be
changed to "exclude_network_and_protocol_config" for IDP SVMDR.
* The property "backoff_level" is only applicable for FlexVol SnapMirror
relationships.

== Examples

The following examples show how to perform the SnapMirror "resync",
"initialize", "resume", "quiesce", and "break" operations. In addition, a
relationship can be failed over to the destination endpoint and start
serving data from the destination endpoint. A failed over relationship can
be failed back to the source endpoint and serve data from the source
endpoint. Also a relationship can be reversed by making the source
endpoint as the new destination endpoint and the destination endpoint as
the new source endpoint.

To update an associated SnapMirror policy.
```

```
----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"policy": { "name" : "MirrorAndVaultDiscardNetwork"}}'
----

To perform SnapMirror "resync" for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored"}'
----

To perform SnapMirror "initialize" for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored"}'
----

To perform SnapMirror "resume" for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored"}'
----

To perform SnapMirror "quiesce" for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff" '{"state":"paused"}'
----

To perform SnapMirror "break" for an asynchronous SnapMirror relationship.
This operation does a failover to the destination endpoint. After a the
failover, data can then be served from the destination endpoint.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff" '{"state":"broken_off"}'
----

To forcefully failover to the destination endpoint and start serving data
```

```
from the destination endpoint.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/?force=true" '{"state":"broken_off"}'
----

To failback to the source endpoint and start serving data from the source
endpoint for an asynchronous relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/?failback=true" '{"state":"snapmirrored"}'
----

To failback to the source endpoint and start serving data from the source
endpoint for a synchronous relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/?failback=true" '{"state":"in_sync"}'
----

To reverse the direction of an asynchronous relationship, that is, make
the source endpoint as the new destination endpoint and make the
destination endpoint as the new source endpoint.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"source": {"path": "dst_svm:dst_vol"}, "destination":
{"path": "src_svm:src_vol"}, "state": "snapmirrored"}'
----

To reverse the direction of a synchronous relationship, that is, make the
source endpoint as the new destination endpoint and make the destination
endpoint as the new source endpoint.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"source": {"path": "dst_svm:dst_vol"}, "destination":
{"path": "src_svm:src_vol"}, "state": "in_sync"}'
----

Updating SnapMirror transfer_schedule and throttle for an asynchronous
SnapMirror relationship. Transfer_schedule can be specified as UUID or
name or both.
```

```
---
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"transfer_schedule":{"uuid":"817500fa-092d-44c5-9c10-
7b54f7b2f20a", "name":"5min"}, "throttle":100}'
----

Removing the SnapMirror transfer_schedule for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"transfer_schedule":{"uuid":null, "name":null}}'
----

Removing the SnapMirror throttle for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"throttle":0}'
----

To perform SnapMirror "resync" and update the SnapMirror transfer_schedule
for an asynchronous SnapMirror relationship. First the transfer_schedule
is modified and then the resync transfer is initiated.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored",
transfer_schedule":{"uuid":"817500fa-092d-44c5-9c10-7b54f7b2f20a",
"name":"5min"}}'
----

To perform SnapMirror "initialize" and update the SnapMirror throttle for
an asynchronous SnapMirror relationship. First the throttle is modified
and then the initialize transfer is initiated. The initialize transfer
will use this new throttle.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored", "throttle":100}'
----

To perform SnapMirror "resync" and update the SnapMirror throttle for an
asynchronous SnapMirror relationship. First the throttle is modified and
then the resync transfer is initiated. The resync transfer will use this
```

```
new throttle.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored", "throttle":100}'
----

To perform a SnapMirror active sync or Asynchronous Consistency Group
expansion.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{ "source" :
{"consistency_group_volumes":[{"name":"vol"}]}, "destination" :
{"consistency_group_volumes":[{"name":"voldp"}]} }'
----

Updating SnapMirror backoff_level for an asynchronous SnapMirror
relationship.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"backoff_level": "none"}'
----

To perform SnapMirror "initialize" and update the SnapMirror backoff_level
for an asynchronous SnapMirror relationship. First the backoff_level is
modified and then the initialize transfer is initiated. The initialize
transfer will use this new backoff_level.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored", "backoff_level": "medium"}'
----

To perform SnapMirror "resync" and update the SnapMirror backoff_level for
an asynchronous SnapMirror relationship. First the backoff_level is
modified and then the resync transfer is initiated. The resync transfer
will use this new backoff_level.

----
 PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-
005056a707ff/" '{"state":"snapmirrored", "backoff_level": "medium"}'
----

== Learn more
```

```
* xref:{relative_path}snapmirror_relationships_endpoint_overview.html[DOC
/snapmirror/relationships]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|SnapMirror relationship UUID


|failover
|boolean
|query
|False
a|If this parameter is set, validation and failover will occur to the SVM-
DR SnapMirror relationship destination endpoint. Any other fields
specified with this parameter will be ignored. This parameter is supported
only for SVM-DR SnapMirror relationships.

* Introduced in: 9.7
* Default value:


|force-failover
|boolean
|query
|False
a|If this parameter is set, failover will occur to the SVM-DR SnapMirror
relationship destination endpoint, overriding the validation errors. Any
other fields specified with this parameter will be ignored. This parameter
is supported only for SVM-DR SnapMirror relationships.

* Introduced in: 9.7
* Default value:
```

```
|force
|boolean
|query
|False
a|If this parameter is set while specifying the state as "broken_off",
indicates a forced failover overriding the validation errors.

* Introduced in: 9.8
* Default value:


|failback
|boolean
|query
|False
a|If this parameter is set while specifying the state as "snapmirrored",
indicates recovery of the failed over SnapMirror relationship by
preserving the data written on the destination endpoint when the
SnapMirror relationship was in failed over state. This flag is only
applicable to SVM-DR SnapMirror relationships.

* Introduced in: 9.8
* Default value:


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|validate_only
|boolean
|query
```

```
|False
a|Validate the operation and its parameters, without actually performing
the operation.

* Introduced in: 9.7

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
a|Snapshot exported to clients on destination.


|group_type
```

```
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
```

```
when the snapshot is restored from a destination endpoint to another
endpoint.


|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
PATCH state is being changed to "broken_off".


|restore_to_snapshot
```

```
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
"sync". A PATCH call on the state change only triggers the transition to
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
```

```
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.


|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.


|uuid
|string
a|Unique identifier of the SnapMirror relationship.
```

```
|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "backoff_level": "medium",
  "destination": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "luns": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "group_type": "consistency_group",
  "identity_preservation": "string",
  "io_serving_copy": "C1_sti85-vsim-ucs209a_cluster, C1_sti85-vsim-
ucs209c_cluster",
  "lag_time": "PT8H35M42S",
  "last_transfer_network_compression_ratio": 61,
  "last_transfer_type": "initialize",
  "master_bias_activated_site": "C1_sti85-vsim-ucs209a_cluster",
  "policy": {
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "preferred_site": "C1_sti85-vsim-ucs209a_cluster",
  "restore": true,
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "name": "cluster1",
```

```
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "luns": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "snapmirrored",
  "throttle": 0,
  "total_transfer_bytes": 1098210312,
  "total_transfer_duration": "PT3M21S",
  "transfer": {
    "end_time": "2020-12-02 18:36:19 -0800",
    "last_updated_time": "2023-09-14 15:39:19 -0700",
    "state": "string",
    "total_duration": "PT28M41S",
    "type": "initialize",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unhealthy_reason": [
    {
      "arguments": [],
      "code": "6621444",
      "message": "Failed to complete update operation on one or more item
relationships."
    },
    {
      "arguments": [],
      "code": "6621445",
      "message": "Group Update failed"
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303825
| Could not retrieve information for the SnapMirror policy type

| 13303817
| Unknown value for the SnapMirror state
```

```
| 13303829
| Invalid state

| 13303830
| Transient state

| 13303831
| Invalid state for async SnapMirror relationship

| 13303834
| Given input valid only for FlexGroup SnapMirror relationship

| 13303835
| Given flag is valid only when PATCH state is broken_off

| 13303836
| Given flag is valid only when PATCH state is snapmirrored or in_sync

| 13303818
| Invalid state transition requested

| 13303828
| Given state change is not possible for SVM SnapMirror relationship

| 13303833
| Requested state change is not possible

| 13303832
| SnapMirror relationship is already initialized

| 13303824
| Quiescing the SnapMirror relationship has failed

| 13303826
| Required environment variables are not set

| 13303827
| Internal Error

| 13303823
| Quiesce operation timed out

| 13303821
| Invalid SnapMirror policy name/UUID

| 13303819
```

```
| Could not retrieve SnapMirror policy information

| 13303851
| Cannot modify attributes of SnapMirror restore relationship

| 13303816
| Could not retrieve the state or status values for the relationship. Wait
a few minutes and try the operation again

| 13303837
| Given flags are valid only if SnapMirror state change is requested

| 6619546
| Destination must be a dp volume

| 13303808
| Transition to broken_off state failed

| 13303809
| Transition to paused state failed

| 13303810
| Transition to snapmirrored state failed

| 13303811
| Transition from paused state failed

| 13303820
| SnapMirror policy, transfer_schedule, and throttle, if specified were
successfully updated, state transition failed

| 13303856
| SVM is not configured with any data protocol

| 13303857
| SVM is not configured with any network interface

| 13303858
| Internal error. Failed to check LIF and protocols details for SVM

| 13303859
| Internal error. SVM Failover operation failed. SVM operational state is
unavailable.

| 13303865
| Modifying the specified SnapMirror policy is not supported.
```

```
| 13303866
| Cannot use the specified policy to modify the policy of the
relationship.

| 13303867
| Modifying the policy of an async-mirror or a vault relationship is not
supported.

| 13303884
| LIF and protocols details are configured incorrectly for SVM.

| 13303996
| The source and destination clusters both have a policy with the same
name, but they have different properties.

| 13304062
| Cannot reverse the direction of a SnapMirror DP relationship when the
source cluster version is earlier than the destination cluster version.

| 13304070
| Remote peer cluster requires the dp_rest_support capability to support
reversing the direction of a DP relationship.

| 13304071
| Failed to access capabilities on remote cluster.

| 13304080
| Specified uuid and name do not match.

| 13304081
| Modifying a property during the operation is not supported.

| 13304082
| The specified properties are mutually exclusive.

| 13304083
| The specified property is not supported because all nodes in the cluster
are not capable of supporting the property.

| 13304086
| Reversing the direction of a SnapMirror relationship associated with a
policy containing the property create_snapshot_on_source set to false is
not supported.

| 6619715
| Modification of relationship is in progress. Retry the command after a
few minutes.
```

```
| 6619699
| Schedule not found.

| 13304108
| Schedule not found in the Administrative SVM or the SVM for the
relationship.

| 13304111
| The SnapMirror active sync relationship consistency groups are nested.
Expanding an SnapMirror active sync relationship with a pre-existing DP
volume is only supported for flat consistency groups.

| 6621458
| The destination Consistency Group is the source of a SnapMirror
Synchronous (SM-S) relationship. Sources of SM-S relationships cannot be
the destination of any other SnapMirror relationship.

| 6621782
| A property of the policy is not valid for relationships between these
endpoints.
|===

| 13304120    | Values specified for the source.path and destination.path
properties do not match the relationship's source.path or destination.path
properties.
| 13304003    | IntraCluster flip operation is not supported. |
| 6621125     | The policy is not valid for relationships with FlexGroup
volume endpoints. Only policies without snapshot creation schedules are
supported for these relationships. |
| 13304093    | The property specified is not supported for the specified
relationships. |
| 6622077     | The expand operation has failed on the SnapMirror active
sync relationship with specified destination path. |

[cols=2*]
|===
| 6619720
| Relationship information has been updated and is being propagated. Wait
a few minutes and try the operation again.

| 13304156
| Either use state as broken-off or failover as true for SVM-DR fail-over
operation.

| 13304159
```

```
| Could not retrieve the state or status values for the relationship. Wait
a few minutes and try the operation again. If the error persists, address
the error condition shown.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

SnapMirror Consistency Group failover error message.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#status]
[.api-collapsible-fifth-title]
status

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|code
|string
a|Status code


|message
|string
a|SnapMirror Consistency Group failover status.


|===


[#snapmirror_consistency_group_failover]
[.api-collapsible-fifth-title]
snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror
Consistency Group failover can be a planned or an unplanned operation.
Only active SnapMirror Consistency Group failover operation progress can
be monitored using this object. In case of an error during the failover
operation, the property "consistency_group_failover.error" holds the
reason for the error. ONTAP automatically retries any failed SnapMirror
Consistency Group failover operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|SnapMirror Consistency Group failover state.


|status
|link:#status[status]
a|

|type
|string
a|SnapMirror Consistency Group failover type.
```

```
|===


[#bucket_retention]
[.api-collapsible-fifth-title]
bucket_retention

Specifies the retention-mode and default retention period configured on
the destination bucket.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|default_period
|string
a|Specifies the default retention period that is applied to objects while
committing them to the WORM state without an associated retention period.
The retention period can be in years, or days. The retention period value
represents a duration and must be specified in the ISO-8601 duration
format.  A period specified for years and days is represented in the ISO-
8601 format as "P+++<num>+++Y" and "P+++<num>+++D" respectively, for
example "P10Y" represents a duration of 10 years. The period string must
contain only a single time element that is, either years, or days. A
duration which combines different periods is not supported, for example
"P1Y10D" is not supported.+++</num>++++++</num>+++


|===


[#storage_service]
[.api-collapsible-fifth-title]
storage_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
```

```
|boolean
a|This property indicates whether to create the destination endpoint using
storage service. This property is supported for Unified ONTAP destination
endpoints only.

* Introduced in: 9.7
* x-nullable: true


|enforce_performance
|boolean
a|Optional property to enforce storage service performance on the
destination endpoint. This property is applicable to FlexVol volume,
FlexGroup volume, and Consistency Group endpoints. This property is
supported for Unified ONTAP destination endpoints only.

* Default value: 1
* Introduced in: 9.7
* x-nullable: true


|name
|string
a|Optional property to specify the storage service name for the
destination endpoint. This property is considered when the property
"create_destination.storage_service.enabled" is set to "true". When the
property "create_destination.storage_service.enabled" is set to "true" and
the "create_destination.storage_service.name" for the endpoint is not
specified, then ONTAP selects the highest storage service available on the
cluster to provision the destination endpoint. This property is applicable
to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This
property is supported for Unified ONTAP destination endpoints only.

* enum: ["extreme", "performance", "value"]
* Introduced in: 9.6
* x-nullable: true


|===


[#tiering]
[.api-collapsible-fifth-title]
tiering

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|policy
|string
a|Optional property to specify the destination endpoint's tiering policy
when "create_destination.tiering.supported" is set to "true". This
property is applicable to FlexVol volume, FlexGroup volume, and
Consistency Group endpoints. This property determines whether the user
data blocks of the destination endpoint in a FabricPool will be tiered to
the cloud store when they become cold. FabricPool combines flash
(performance tier) with a cloud store into a single aggregate. Temperature
of the destination endpoint volume blocks increases if they are accessed
frequently and decreases when they are not.
all &dash; This policy allows tiering of both destination endpoint
snapshots and the user transferred data blocks to the cloud store as soon
as possible by ignoring the temperature on the volume blocks. This tiering
policy is not applicable for Consistency Group destination endpoints or
for synchronous relationships.
auto &dash; This policy allows tiering of both destination endpoint
snapshots and the active file system user data to the cloud store
none &dash; Destination endpoint volume blocks will not be tiered to the
cloud store.
snapshot_only &dash; This policy allows tiering of only the destination
endpoint volume snapshots not associated with the active file system. The
default tiering policy is "snapshot_only" for a FlexVol volume and "none"
for a FlexGroup volume. This property is supported for Unified ONTAP
destination endpoints only.

* enum: ["all", "auto", "none", "snapshot_only"]
* Introduced in: 9.6
* x-nullable: true


|supported
|boolean
a|Optional property to enable provisioning of the destination endpoint
volumes on FabricPool aggregates. This property is applicable to FlexVol
volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool
aggregates are used if this property is set to "true" and only non
FabricPool aggregates are used if this property is set to "false". Tiering
support for a FlexGroup volume can be changed by moving all of the
constituents to the required aggregates. Note that in order to tier data,
not only do the destination endpoint volumes need to support tiering by
using FabricPools, the "create_destination.tiering.policy" must not be
```

```
"none". A destination endpoint that uses FabricPools but has a tiering
"policy" of "none" supports tiering but will not tier any data. This
property is supported for Unified ONTAP destination endpoints only.

* Introduced in: 9.6
* x-nullable: true


|===


[#snapmirror_destination_creation]
[.api-collapsible-fifth-title]
snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a
SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM,
Consistency Group or ONTAP S3 Bucket. Given a source endpoint, the
destination endpoint is provisioned in the SVM specified in the
"destination.path" property. While protecting an SVM, the SVM destination
endpoint can only be provisioned on the local cluster. To provision the
SVM destination endpoint use the optional "source.cluster.name" property
to specify the remote cluster name or use the optional
"source.cluster.uuid" property to specify the remote cluster UUID. When
"create_destination.enabled" option is specified while making a POST for a
SnapMirror relationship, the relationship can be automatically initialized
by setting the "state" either to "snapmirrored" when the policy is of type
"async" or to "in_sync" when the policy is of type "sync". The
"destination.path" property must specify the destination endpoint path.
For example, for FlexVol volume and FlexGroup volume, the
"destination.path" can be specified as +++<destination-svm-name:dp-volume-
name>+++, for SVM data protection, the "destination.path" must be
specified as <destination-svm-name:>, and for Consistency Group, the
"destination.path" must be specified as <destination-svm-name:> along with
the "destination.consistency_group_volumes" or "destination.luns" property
to indicate the list of destination volumes or LUNs of type "DP" in the
Consistency Group. For a FlexVol volume, a FlexGroup volume, Consistency
Group or a Bucket destination endpoint, the properties in this object can
be specified either from the source or the destination cluster. For an SVM
destination endpoint, the properties in this object can be specified from
the destination cluster only. This object is not supported for non ONTAP
endpoints. While protecting a S3 Bucket, the optional "size" property can
be used to create ONTAP S3 Bucket destination endpoint of the specified
size.</destination-svm-name:></destination-svm-name:>+++</destination-svm-
name:dp-volume-name>+++
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|bucket_retention
|link:#bucket_retention[bucket_retention]
a|Specifies the retention-mode and default retention period configured on
the destination bucket.


|enabled
|boolean
a|Optional property to create the destination endpoint when establishing a
SnapMirror relationship. It is assumed to be "false" if no other property
is set and assumed to be "true" if any other property is set.


|size
|integer
a|Optional property to specify the size of destination endpoint in bytes.
This property is applicable only to ONTAP S3 Bucket endpoints. The minimum
size for S3 bucket is 80MB and maximum size is 64TB. If not specified,
system will create destination with default size of 800GB.

* Introduced in: 9.10
* x-nullable: true


|snapshot_locking_enabled
|boolean
a|Optional property to create the destination endpoint with snapshot
locking enabled when establishing a SnapMirror relationship. This property
is applicable to FlexVol volumes and FlexGroup volumes.


|storage_service
|link:#storage_service[storage_service]
a|

|tiering
|link:#tiering[tiering]
a|

|===
```

```
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
```

```
PATCH method.


|===


[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore
```

```
* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#policy]
[.api-collapsible-fifth-title]
policy

Basic policy information of the relationship.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the SnapMirror policy.


|type
|string
a|

|uuid
|string
a|Unique identifier of the SnapMirror policy.


|===


[#snapmirror_source_endpoint]
[.api-collapsible-fifth-title]
snapmirror_source_endpoint
```

```
Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.14
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===
```

```
[#svmdr_volumes]
[.api-collapsible-fifth-title]
svmdr_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#transfer]
[.api-collapsible-fifth-title]
transfer

Basic information on the current transfer or the last transfer if there is
no active transfer at the time of the request.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_transferred
|integer
a|Total bytes transferred in the current or last successful transfer.


|end_time
|string
a|End time of the last transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in the current transfer.
```

```
|state
|string
a|

|total_duration
|string
a|Elapsed time to transfer all snapshots for the last successful transfer.


|type
|string
a|Specifies the operation type of the current transfer on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from "uninitialized" to "snapmirrored" or "in_sync". The _update_
transfer occurs when snapshots are being transferred from the source
endpoint to the destination endpoint as part of a scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from "broken_off" to "snapmirrored" or "in_sync". The _restore_ transfer
occurs when a snapshot is being restored from a destination endpoint to
another endpoint.


|uuid
|string
a|Transfer UUID. This property is applicable only for active transfers.


|===


[#transfer_schedule]
[.api-collapsible-fifth-title]
transfer_schedule

Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapmirror_error]
[.api-collapsible-fifth-title]
snapmirror_error

SnapMirror error


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the error message encountered.


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#snapmirror_relationship]
[.api-collapsible-fifth-title]
snapmirror_relationship
```

```
SnapMirror relationship information. The SnapMirror relationship can be
either "async" or "sync" based on the type of SnapMirror policy associated
with the relationship. The source and destination endpoints of a
SnapMirror relationship must be of the same type, for example, if the
source endpoint is a FlexVol volume then the destination endpoint must be
a FlexVol volume.
The SnapMirror policy type "async" can be used when the SnapMirror
relationship has FlexVol volume or FlexGroup volume or SVM as the
endpoint. The SnapMirror policy type "sync" can be used when the
SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror
policy type "sync" with "sync_type" as "automated_failover" can be used
when the SnapMirror relationship has Consistency Group as the endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|backoff_level
|string
a|Specifies the SnapMirror backoff level due to Client Ops for FlexVol
SnapMirror relationships.


|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|exported_snapshot
|string
```

```
a|Snapshot exported to clients on destination.


|group_type
|string
a|Specifies the group type of the top level SnapMirror relationship. The
volume relationships are shown as _none_, the SVMDR relationships are
shown as _svm_dr_, the Consistency Group relationships are shown as
_consistency_group_, and the FlexGroup volume relationships are shown as
_flexgroup_.


|healthy
|boolean
a|Is the relationship healthy?


|identity_preservation
|string
a|Specifies which configuration of the source SVM is replicated to the
destination SVM. This property is applicable only for SVM data protection
with "async" policy type. This "identity_preservation" overrides the
"identity_preservation" set on the SnapMirror relationship's policy.


|io_serving_copy
|string
a|Specifies the sites serving I/O for the SnapMirror active sync
relationship.


|lag_time
|string
a|Time since the exported snapshot was created.


|last_transfer_network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled for the last successful transfer.


|last_transfer_type
|string
a|Specifies the operation type of the last transfer that occurred on the
relationship. The _initialize_ transfer occurs when the relationship state
changes from uninitialized to snapmirrored or in_sync. The _update_
```

```
transfer occurs when the snapshots are transferred from the source
endpoint to the destination endpoint as part of scheduled or manual
update. The _resync_ transfer occurs when the relationship state changes
from broken_off to snapmirrored or in_sync. The _restore_ transfer occurs
when the snapshot is restored from a destination endpoint to another
endpoint.


|master_bias_activated_site
|string
a|Specifies the Master Bias Activated Site for the SnapMirror active sync
relationship.


|policy
|link:#policy[policy]
a|Basic policy information of the relationship.


|preferred_site
|string
a|Specifies the Primary Site of the SnapMirror active sync relationship.


|preserve
|boolean
a|Set to true on resync to preserve snapshots on the destination that are
newer than the latest common snapshot. This property is applicable only
for relationships with FlexVol volume or FlexGroup volume endpoints and
when the PATCH state is being changed to "snapmirrored".


|quick_resync
|boolean
a|Set to true to reduce resync time by not preserving storage efficiency.
This property is applicable only for relationships with FlexVol volume
endpoints and SVMDR relationships when the PATCH state is being changed to
"snapmirrored".


|recover_after_break
|boolean
a|Set to true to recover from a failed SnapMirror break operation on a
FlexGroup volume relationship. This restores all destination FlexGroup
constituent volumes to the latest snapshot, and any writes to the read-
write constituents are lost. This property is applicable only for
SnapMirror relationships with FlexGroup volume endpoints and when the
```

```
PATCH state is being changed to "broken_off".


|restore_to_snapshot
|string
a|Specifies the snapshot to restore to on the destination during the break
operation. This property is applicable only for SnapMirror relationships
with FlexVol volume endpoints and when the PATCH state is being changed to
"broken_off".


|source
|link:#snapmirror_source_endpoint[snapmirror_source_endpoint]
a|Source endpoint of a SnapMirror relationship. For a GET request, the
property "cluster" is populated when the endpoint is on a remote cluster.
A POST request to establish a SnapMirror relationship between the source
endpoint and destination endpoint and when the source SVM and the
destination SVM are not peered, must specify the "cluster" property for
the remote endpoint.


|state
|string
a|State of the relationship.
To initialize the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync".
To break the relationship, PATCH the state to "broken_off" for
relationships with a policy of type "async" or "sync". SnapMirror
relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "broken_off".
To resync the relationship, PATCH the state to "snapmirrored" for
relationships with a policy of type "async" or to state "in_sync" for
relationships with a policy of type "sync". SnapMirror relationships with
the policy type as "sync" and "sync_type" as "automated_failover" can be
in "broken_off" state due to a failed attempt of SnapMirror failover.
To pause the relationship, suspending further transfers, PATCH the state
to "paused" for relationships with a policy of type "async" or "sync".
SnapMirror relationships with the policy type as "sync" and "sync_type" as
"automated_failover" cannot be "paused".
To resume transfers for a paused relationship, PATCH the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".
The entries "in_sync", "out_of_sync", "synchronizing", "expanding", and
"shrinking" are only applicable to relationships with a policy of type
"sync". A PATCH call on the state change only triggers the transition to
```

```
the specified state. You must poll on the "state", "healthy" and
"unhealthy_reason" properties using a GET request to determine if the
transition is successful. To automatically initialize the relationship
when specifying "create_destination" property, set the state to
"snapmirrored" for relationships with a policy of type "async" or to state
"in_sync" for relationships with a policy of type "sync".


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship's policy. If neither of these are set,
defaults to 0, which is interpreted as unlimited.


|total_transfer_bytes
|integer
a|Cumulative bytes transferred for the relationship.


|total_transfer_duration
|string
a|Indicates the cumulative duration of all transfers since the last
aggregate relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that hosts the relationship.


|transfer
|link:#transfer[transfer]
a|Basic information on the current transfer or the last transfer if there
is no active transfer at the time of the request.


|transfer_schedule
|link:#transfer_schedule[transfer_schedule]
a|Schedule used to update asynchronous relationships. This
"transfer_schedule" overrides the "transfer_schedule" set on the
SnapMirror relationship's policy. To remove the "transfer_schedule", set
its value to null (no-quotes). Only cron schedules are supported for
SnapMirror.


|unhealthy_reason
|array[link:#snapmirror_error[snapmirror_error]]
a|Reason the relationship is not healthy. It is a concatenation of up to
four levels of error messages.
```

```
|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage SnapMirror relationship tranfers

:leveloffset: +1


[[ID25973d210d66023059282f2c63ce1359]]
= Manage SnapMirror relationship tranfers



== Overview

This API is used to manage transfers on an existing SnapMirror
relationship.

You can initiate SnapMirror operations such as "initialize", "update",
"restore-transfer", and "abort" using this API on asynchronous SnapMirror
relationship. On a synchronous SnapMirror relationship, you can initiate
SnapMirror "initialize" operation. The GET for this API reports the status
of both active transfers and transfers that have terminated within the
past 24 hours.
For the restore relationships, the POST on transfers API triggers
"restore-transfer". Successful completion of "restore" also deletes the
restore relationship. If the "restore" fails, DELETE on relationships must
be called to delete the restore relationship.
```

```
A transfer on an asynchronous SnapMirror relationship with Application
Consistency Group endpoints expands the destination Application
Consistency Group endpoint if the source Application Consistency Group
endpoint is already expanded.





[[ID8a9e44a95accfeb189aed26d7308c3fd]]
= Retrieve ongoing SnapMirror transfers for a relationship

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/relationships/{relationship.uuid}/transfers`#

*Introduced In:* 9.6

Retrieves the list of ongoing SnapMirror transfers for the specified
relationship.

== Related ONTAP commands

* `snapmirror show`

== Example

----
GET "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-
005056a793dd/transfers"
----

== Learn more

*
xref:{relative_path}snapmirror_relationships_relationship.uuid_transfers_e
ndpoint_overview.html[DOC
/snapmirror/relationships/{relationship.uuid}/transfers]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|relationship.uuid
|string
|path
|True
a|SnapMirror relationship UUID


|snapshot
|string
|query
|False
a|Filter by snapshot


|error_info.message
|string
|query
|False
a|Filter by error_info.message

* Introduced in: 9.10


|error_info.code
|integer
|query
|False
a|Filter by error_info.code

* Introduced in: 9.10


|bytes_transferred
|integer
|query
|False
a|Filter by bytes_transferred


|uuid
|string
|query
|False
```

```
a|Filter by uuid


|relationship.restore
|boolean
|query
|False
a|Filter by relationship.restore


|relationship.destination.luns.uuid
|string
|query
|False
a|Filter by relationship.destination.luns.uuid

* Introduced in: 9.16


|relationship.destination.luns.name
|string
|query
|False
a|Filter by relationship.destination.luns.name

* Introduced in: 9.16


|relationship.destination.svm.name
|string
|query
|False
a|Filter by relationship.destination.svm.name


|relationship.destination.svm.uuid
|string
|query
|False
a|Filter by relationship.destination.svm.uuid


|relationship.destination.consistency_group_volumes.name
|string
|query
|False
a|Filter by relationship.destination.consistency_group_volumes.name
```

```
* Introduced in: 9.8


|relationship.destination.cluster.name
|string
|query
|False
a|Filter by relationship.destination.cluster.name

* Introduced in: 9.7


|relationship.destination.cluster.uuid
|string
|query
|False
a|Filter by relationship.destination.cluster.uuid

* Introduced in: 9.7


|relationship.destination.path
|string
|query
|False
a|Filter by relationship.destination.path


|state
|string
|query
|False
a|Filter by state


|last_updated_time
|string
|query
|False
a|Filter by last_updated_time

* Introduced in: 9.14


|on_demand_attrs
|string
```

```
|query
|False
a|Filter by on_demand_attrs

* Introduced in: 9.13


|total_duration
|string
|query
|False
a|Filter by total_duration

* Introduced in: 9.10


|checkpoint_size
|integer
|query
|False
a|Filter by checkpoint_size


|end_time
|string
|query
|False
a|Filter by end_time

* Introduced in: 9.10


|network_compression_ratio
|string
|query
|False
a|Filter by network_compression_ratio

* Introduced in: 9.13


|throttle
|integer
|query
|False
a|Filter by throttle
```

```
* Introduced in: 9.9


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snapmirror_transfer[snapmirror_transfer]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
```

```
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "bytes_transferred": 0,
      "checkpoint_size": 0,
      "end_time": "2020-12-02 18:36:19 -0800",
      "error_info": {
        "code": 6620046,
        "message": "Transfer aborted"
      },
      "last_updated_time": "2023-09-15 16:58:39 -0700",
      "network_compression_ratio": 61,
      "on_demand_attrs": "read_write_with_user_data_pull",
      "relationship": {
        "destination": {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "consistency_group_volumes": [
            {
              "name": "volume1"
            }
          ],
          "luns": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "/vol/volume1/lun1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "path": "svm1:volume1",
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
```

```
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          }
        },
        "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
      },
      "snapshot": "string",
      "state": "string",
      "storage_efficiency_enabled": true,
      "throttle": 0,
      "total_duration": "PT28M41S",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
```

```
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_info]
[.api-collapsible-fifth-title]
error_info

Error information for the transfer.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Error code


|message
|string
a|Error message


|===
```

```
[#files]
[.api-collapsible-fifth-title]
files

Specifies a file or LUN consisting of a source_path and an optional
destination_path. If not specified, the destination_path is the same as
the source_path. File restore is not supported if the source_path or
destination_path contains commas in its directory or file name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination_path
|string
a|

|source_path
|string
a|

|===


[#options]
[.api-collapsible-fifth-title]
options
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
```

```
|string
a|

|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
```

```
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#relationship]
[.api-collapsible-fifth-title]
relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|restore
|boolean
a|Is the relationship for restore?
```

```
|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#snapmirror_transfer]
[.api-collapsible-fifth-title]
snapmirror_transfer

SnapMirror transfer information


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|bytes_transferred
|integer
a|Bytes transferred


|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
|link:#error_info[error_info]
a|Error information for the transfer.
```

```
|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.


|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|relationship
|link:#relationship[relationship]
a|

|snapshot
|string
a|Name of snapshot being transferred.


|state
|string
a|Status of the transfer. Set PATCH state to "aborted" to abort the
transfer. Set PATCH state to "hard_aborted" to abort the transfer and
discard the restart checkpoint. To find "queued" transfers refer to
relationships GET API.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship or SnapMirror relationship's policy. If
neither of these are set, defaults to 0, which is interpreted as
unlimited.
```

```
|total_duration
|string
a|Elapsed transfer time.


|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID28662fde77b5b962224e5d4062007de5]]
= Start a SnapMirror transfer operation

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/snapmirror/relationships/{relationship.uuid}/transfers`#

*Introduced In:* 9.6

Starts a SnapMirror transfer operation. This API initiates a restore
operation if the SnapMirror relationship is of type "restore". Otherwise,
it initiates a SnapMirror "initialize" operation or "update" operation
based on the current SnapMirror state.

== Default property values

* `storage_efficiency_enabled` - _true_

== Related ONTAP commands

* `snapmirror update`
```

```
* `snapmirror initialize`
* `snapmirror restore`

== Examples

The following examples show how to perform SnapMirror "initialize",
"update", and "restore" operations.

Perform SnapMirror initialize or update

----
 POST "/api/snapmirror/relationships/e4e7e130-0279-11e9-b566-
0050568e9909/transfers" '{}'
----

Perform SnapMirror initialize, update or restore with throttle value set

----
 POST "/api/snapmirror/relationships/e4e7e130-0279-11e9-b566-
0050568e9909/transfers" '{"throttle":"100"}'
----

Perform SnapMirror restore transfer of a file

----
 POST "/api/snapmirror/relationships/c8c62a90-0fef-11e9-b09e-
0050568e7067/transfers" '{"source_snapshot": "src",
"files":[{"source_path": "/a1.txt.0", "destination_path": "/a1-
renamed.txt.0"}]}'
----

Performing a SnapMirror initialize or update using a particular snapshot.

----
 POST "/api/snapmirror/relationships/e4e7e130-0279-11e9-b566-
0050568e9909/transfers" '{"source_snapshot":"snap1"}'
----

== Learn more

*
xref:{relative_path}snapmirror_relationships_relationship.uuid_transfers_e
ndpoint_overview.html[DOC
/snapmirror/relationships/{relationship.uuid}/transfers]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|relationship.uuid
|string
|path
|True
a|SnapMirror relationship UUID


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_transferred
|integer
a|Bytes transferred


|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
|link:#error_info[error_info]
a|Error information for the transfer.


|files
|array[link:#files[files]]
a|This is supported for transfer of restore relationship only. This
specifies the list of files or LUNs to be restored. Can contain up to
eight files or LUNs.


|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.
```

```
|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|options
|array[link:#options[options]]
a|Options for snapmirror transfer.


|relationship
|link:#relationship[relationship]
a|

|snapshot
|string
a|Name of snapshot being transferred.


|source_snapshot
|string
a|Specifies the snapshot on the source to be transferred to the
destination.


|storage_efficiency_enabled
|boolean
a|This is supported for transfer of restore relationship only. Set this
property to "false" to turn off storage efficiency for data transferred
over the wire and written to the destination.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship or SnapMirror relationship's policy. If
neither of these are set, defaults to 0, which is interpreted as
unlimited.


|total_duration
|string
a|Elapsed transfer time.
```

```
|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "end_time": "2020-12-02 18:36:19 -0800",
  "error_info": {
    "code": 6620046,
    "message": "Transfer aborted"
  },
  "files": [
    {
      "destination_path": "/dirb/file2",
      "source_path": "/dira/file1"
    }
  ],
  "last_updated_time": "2023-09-15 16:58:39 -0700",
  "network_compression_ratio": 61,
  "on_demand_attrs": "read_write_with_user_data_pull",
  "options": [
    {}
  ],
  "relationship": {
    "destination": {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "consistency_group_volumes": [
        {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ],
```

```
      "ipspace": "Default",
      "luns": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "path": "svm1:volume1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "source_snapshot": "string",
  "throttle": 0,
  "total_duration": "PT28M41S",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====
```


```
== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303845
| Restore operation failed

| 13303812
| Initialize operation failed

| 13303844
| Update operation failed

| 13303846
| Empty source path file list

| 13303847
| Invalid property

| 13304040
| Throttle not supported for update of Synchronous SnapMirror
relationships

| 6620237
| SnapMirror relationship database write failed

| 6620238
| SnapMirror relationship database read failed
|===

| 13304134    | The overwrite property specified does not match the value
specified in the previous request. ||
| 6619720     | Relationship information has been updated and is being
propagated. Wait a few minutes and try the operation again. |
| 13303816    | Could not retrieve the state or status values for the
relationship. Wait a few minutes and try the operation again. |
| 13304159    | Could not retrieve the state or status values for the
relationship. Wait a few minutes and try the operation again. If the error
persists, address the error condition shown. |


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#error_info]
[.api-collapsible-fifth-title]
error_info

Error information for the transfer.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Error code


|message
|string
a|Error message
```

```
|===


[#files]
[.api-collapsible-fifth-title]
files

Specifies a file or LUN consisting of a source_path and an optional
destination_path. If not specified, the destination_path is the same as
the source_path. File restore is not supported if the source_path or
destination_path contains commas in its directory or file name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination_path
|string
a|

|source_path
|string
a|

|===


[#options]
[.api-collapsible-fifth-title]
options
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|
```

```
|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|uuid
|string
a|Unique identifier of the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* x-ntap-createOnly: true
* Introduced in: 9.8
* x-nullable: true


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|ipspace
|string
a|Optional property to specify the IPSpace of the SVM.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
```

```
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#relationship]
[.api-collapsible-fifth-title]
relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.
```

```
* Introduced in: 9.6


|restore
|boolean
a|Is the relationship for restore?


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#snapmirror_transfer]
[.api-collapsible-fifth-title]
snapmirror_transfer

SnapMirror transfer information


[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_transferred
|integer
a|Bytes transferred


|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
```

```
|link:#error_info[error_info]
a|Error information for the transfer.


|files
|array[link:#files[files]]
a|This is supported for transfer of restore relationship only. This
specifies the list of files or LUNs to be restored. Can contain up to
eight files or LUNs.


|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.


|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|options
|array[link:#options[options]]
a|Options for snapmirror transfer.


|relationship
|link:#relationship[relationship]
a|

|snapshot
|string
a|Name of snapshot being transferred.


|source_snapshot
```

```
|string
a|Specifies the snapshot on the source to be transferred to the
destination.


|storage_efficiency_enabled
|boolean
a|This is supported for transfer of restore relationship only. Set this
property to "false" to turn off storage efficiency for data transferred
over the wire and written to the destination.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship or SnapMirror relationship's policy. If
neither of these are set, defaults to 0, which is interpreted as
unlimited.


|total_duration
|string
a|Elapsed transfer time.


|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID0c62cbd40fdc4386aa298141ae4f8f2d]]
= Retrieve attributes of an ongoing SnapMirror transfer

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/snapmirror/relationships/{relationship.uuid}/transfers/{uuid}`#

*Introduced In:* 9.6

Retrieves the attributes of a specific ongoing SnapMirror transfer.

== Related ONTAP commands

* `snapmirror show`

== Example

----
GET "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-
005056a793dd/transfers/293baa53-e63d-11e8-bff1-005056a793dd"
----

== Learn more

*
xref:{relative_path}snapmirror_relationships_relationship.uuid_transfers_e
ndpoint_overview.html[DOC
/snapmirror/relationships/{relationship.uuid}/transfers]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|relationship.uuid
|string
|path
|True
a|SnapMirror relationship UUID
```

```
|uuid
|string
|path
|True
a|SnapMirror transfer UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|bytes_transferred
|integer
a|Bytes transferred


|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
```

```
|link:#error_info[error_info]
a|Error information for the transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.


|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|relationship
|link:#relationship[relationship]
a|

|snapshot
|string
a|Name of snapshot being transferred.


|state
|string
a|Status of the transfer. Set PATCH state to "aborted" to abort the
transfer. Set PATCH state to "hard_aborted" to abort the transfer and
discard the restart checkpoint. To find "queued" transfers refer to
relationships GET API.


|throttle
|integer
a|Throttle, in KBs per second. This "throttle" overrides the "throttle"
set on the SnapMirror relationship or SnapMirror relationship's policy. If
neither of these are set, defaults to 0, which is interpreted as
```

```
unlimited.


|total_duration
|string
a|Elapsed transfer time.


|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "end_time": "2020-12-02 18:36:19 -0800",
  "error_info": {
    "code": 6620046,
    "message": "Transfer aborted"
  },
  "last_updated_time": "2023-09-15 16:58:39 -0700",
  "network_compression_ratio": 61,
  "on_demand_attrs": "read_write_with_user_data_pull",
  "relationship": {
    "destination": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "consistency_group_volumes": [
        {
          "name": "volume1"
        }
      ],
      "luns": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "path": "svm1:volume1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "state": "string",
  "storage_efficiency_enabled": true,
  "throttle": 0,
  "total_duration": "PT28M41S",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_info]
[.api-collapsible-fifth-title]
error_info

Error information for the transfer.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Error code


|message
|string
a|Error message


|===
```

```
[#files]
[.api-collapsible-fifth-title]
files

Specifies a file or LUN consisting of a source_path and an optional
destination_path. If not specified, the destination_path is the same as
the source_path. File restore is not supported if the source_path or
destination_path contains commas in its directory or file name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination_path
|string
a|

|source_path
|string
a|

|===


[#options]
[.api-collapsible-fifth-title]
options
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
```

```
|string
a|

|uuid
|string
a|

|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|consistency_group_volumes
|array[link:#consistency_group_volumes[consistency_group_volumes]]
a|This property specifies the list of FlexVol volumes or LUNs of a
Consistency Group. Optional on the ASA r2 platform. Mandatory for all
other platforms.


|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
```

```
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#relationship]
[.api-collapsible-fifth-title]
relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|restore
|boolean
a|Is the relationship for restore?
```

```
|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID97f756336dd450a031bb25987b8699cd]]
= Cancel an ongoing SnapMirror transfer

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/snapmirror/relationships/{relationship.uuid}/transfers/{uuid}`#

*Introduced In:* 9.6

Aborts an ongoing SnapMirror transfer. This operation is applicable on
asynchronous SnapMirror relationships.

== Related ONTAP commands

* `snapmirror abort`

== Example

----
PATCH "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-
005056a793dd/transfers/293baa53-e63d-11e8-bff1-005056a793dd"
'{"state":"aborted"}'
----
```

```
== Learn more

*
xref:{relative_path}snapmirror_relationships_relationship.uuid_transfers_e
ndpoint_overview.html[DOC
/snapmirror/relationships/{relationship.uuid}/transfers]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|relationship.uuid
|string
|path
|True
a|SnapMirror relationship UUID


|uuid
|string
|path
|True
a|SnapMirror transfer UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_transferred
|integer
a|Bytes transferred
```

```
|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
|link:#error_info[error_info]
a|Error information for the transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.


|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|relationship
|link:#relationship[relationship]
a|

|snapshot
|string
a|Name of snapshot being transferred.
```

```
|state
|string
a|Status of the transfer. Set PATCH state to "aborted" to abort the
transfer. Set PATCH state to "hard_aborted" to abort the transfer and
discard the restart checkpoint. To find "queued" transfers refer to
relationships GET API.


|total_duration
|string
a|Elapsed transfer time.


|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "end_time": "2020-12-02 18:36:19 -0800",
  "error_info": {
    "code": 6620046,
    "message": "Transfer aborted"
  },
  "last_updated_time": "2023-09-15 16:58:39 -0700",
  "network_compression_ratio": 61,
  "on_demand_attrs": "read_write_with_user_data_pull",
  "relationship": {
    "destination": {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "luns": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "path": "svm1:volume1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "state": "string",
  "storage_efficiency_enabled": true,
  "total_duration": "PT28M41S",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====
```


```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response codes

|===
| Error code | Description

| 13303848
| Abort of sync SnapMirror is not allowed

| 13303849
| SnapMirror transfer state is invalid
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#error_info]
[.api-collapsible-fifth-title]
error_info

Error information for the transfer.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|integer
a|Error code


|message
|string
a|Error message


|===


[#files]
```

```
[.api-collapsible-fifth-title]
files

Specifies a file or LUN consisting of a source_path and an optional
destination_path. If not specified, the destination_path is the same as
the source_path. File restore is not supported if the source_path or
destination_path contains commas in its directory or file name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination_path
|string
a|

|source_path
|string
a|

|===


[#options]
[.api-collapsible-fifth-title]
options
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#consistency_group_volumes]
[.api-collapsible-fifth-title]
consistency_group_volumes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume.


|===


[#luns]
[.api-collapsible-fifth-title]
luns

Optional property for a SnapMirror endpoint. Specifies the list of source
LUNs and optionally list of destination LUNs during SnapMirror Consistency
Group LUN restore operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of a LUN.
A LUN is located within a volume. Optionally, it can be located within a
qtree in a volume.

LUN names are paths of the form "/vol/<volume>[/<qtree>]/<namespace>"
where the qtree name is optional.


|uuid
```

```
|string
a|The unique identifier of the LUN.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapmirror_endpoint]
[.api-collapsible-fifth-title]
snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
```

```
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|luns
|link:#luns[luns]
a|Optional property for a SnapMirror endpoint. Specifies the list of
source LUNs and optionally list of destination LUNs during SnapMirror
Consistency Group LUN restore operation.


|path
|string
a|ONTAP FlexVol/FlexGroup - svm1:volume1
ONTAP SVM               - svm1:
ONTAP Consistency Group - svm1:/cg/cg_name
ONTAP S3                - svm1:/bucket/bucket1
NON-ONTAP               - objstore1:/objstore

* example: svm1:volume1
* Introduced in: 9.6
* x-nullable: true


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|===


[#relationship]
[.api-collapsible-fifth-title]
relationship

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|destination
|link:#snapmirror_endpoint[snapmirror_endpoint]
a|Endpoint of a SnapMirror relationship. For a GET request, the property
"cluster" is populated when the endpoint is on a remote cluster. A POST
request to create the destination SVM endpoint or to establish an SVM DR
relationship must have the property "cluster" populated with the remote
cluster details. A POST request to create the destination FlexVol volume,
FlexGroup volume, Consistency Group, ONTAP S3 bucket and NON-ONTAP object-
store endpoints can optionally specify the "cluster" property when the
source SVM and the destination SVM are peered. A POST request to establish
a SnapMirror relationship between the source endpoint and destination
endpoint and when the source SVM and the destination SVM are not peered,
must specify the "cluster" property for the remote endpoint.

* Introduced in: 9.6


|restore
|boolean
a|Is the relationship for restore?


|uuid
|string
a|Unique identifier of the SnapMirror relationship.


|===


[#snapmirror_transfer]
[.api-collapsible-fifth-title]
snapmirror_transfer

SnapMirror transfer information


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|bytes_transferred
|integer
a|Bytes transferred


|checkpoint_size
|integer
a|Amount of data transferred in bytes as recorded in the restart
checkpoint.


|end_time
|string
a|End time of the transfer.


|error_info
|link:#error_info[error_info]
a|Error information for the transfer.


|last_updated_time
|string
a|Last updated time of the bytes transferred in an active transfer.


|network_compression_ratio
|string
a|Specifies the compression ratio achieved for the data sent over the wire
with network compression enabled. This property is only valid for active
transfers.


|on_demand_attrs
|string
a|Specifies whether or not an on-demand restore is being carried out. This
is only supported for the transfer of restore relationships for entire
volumes from the object store. A value for read_write_with_user_data_pull
should be provided to start an on-demand restore. A file restore from the
object store does not support this option.


|relationship
|link:#relationship[relationship]
a|
```

```
|snapshot
|string
a|Name of snapshot being transferred.


|state
|string
a|Status of the transfer. Set PATCH state to "aborted" to abort the
transfer. Set PATCH state to "hard_aborted" to abort the transfer and
discard the restart checkpoint. To find "queued" transfers refer to
relationships GET API.


|total_duration
|string
a|Elapsed transfer time.


|uuid
|string
a|Unique identifier of the SnapMirror transfer.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


:leveloffset: -1


= Storage
```

```
:leveloffset: +1


[[ID928247161a7e0db6db5e609e5ded7c0]]
= ONTAP REST API Storage endpoints



== Overview

The ONTAP storage APIs can be used to manage physical and logical storage.
This includes management of aggregates, volumes, LUNs, qtrees, snapshots,
quotas, and storage efficiency.





= Retrieve or create a collection of storage aggregates

:leveloffset: +1


[[ID55a204066d61f09f0d6ecb382e625cda]]
= Retrieve or create a collection of storage aggregates



== Retrieving storage aggregate information

The Storage Aggregate GET API retrieves all data aggregates in the
cluster. System owned root aggregates are not included in the output.
This API also supports specific queries, in addition to queries on
aggregate body properties, which affect the output of the API. The
parameters
for these queries are "recommend" and "show_spares". Using the "recommend"
query returns the list of aggregates that are
recommended for creation in the cluster. The "show_spares" query returns a
response outside of the records body, which includes the groups
of usable spares in the cluster. The "show_spares" query along with
"flash_pool_eligible" restricts groups of usable spares for whole disk
based flash pool creation. For storage pool creation, the appropriate
groups of eligible spares will be returned by a GET on
storage/pools?show_spares=true.
The usable count for each class of spares does not include reserved spare
capacity recommended by ONTAP best practices.
```

```
The collection GET returns the aggregate identifiers, UUID and name, and
the node on which the aggregate resides. The instance GET, by default,
returns all of the properties defined in the aggregates object, except
advanced properties.
The properties "space.footprint", "is_spare_low",
"space.block_storage.inactive_user_data",
"space.block_storage.inactive_user_data_percent",
"inode_attributes.version", "inode_attributes.files_private_used",
"inode_attributes.file_private_capacity" and
"inode_attributes.file_public_capacity"  are considered advanced
properties and only returned when requested
using the "fields" query parameter. Performance "metric" and "statistics"
for aggregates are also only returned when requested. The "statistics"
property accounts for the cumulative raw values collected by ONTAP for an
aggregate, while the "metric" property displays the incremental average
for latency and incremental changes in IOPs and throughput over the last
15 seconds. Any external application can use the raw statistics to derive
its own incremental performance metrics.
The collection GET does not support queries in conjunction with the
"order_by" parameter.

== Creating storage aggregates

When the POST command is issued with no properties, the system evaluates
the cluster attached storage, determines the optimal aggregate layout and
configures the aggregates. This layout is completely controlled by the
system.
To view the recommended optimal layout rather than creating it, use the
GET endpoint, setting the "recommend" query to 'true'. This response will
also include any warnings related to any inefficiencies in the storage
configuration.
Recommended aggregate creation is not supported on ONTAP Cloud and
MetroCluster with Fibre Channel (FC).
Alternatively, POST can be used with specific properties to create an
aggregate as requested. At a minimum, the
aggregate name, disk count, and the node where it should reside are
required if any properties are provided.
Aggregates are automatically created and expanded in the ASA r2 system.
Therefore, automatic and manual aggregate creation or expansion through
REST is not supported.

When using POST with input properties, three properties are required.
These are:

* name
```

```
* node.name or node.uuid
* block_storage.primary.disk_count

=== Remaining properties are optional

The following properties can be specified in POST:

* name - Name of the aggregate.
* node.name and node.uuid - Node on which the aggregate will be created.
* block_storage.primary.disk_count - Number of disks to be used to create
the aggregate.
* block_storage.mirror.enabled - Specifies whether or not the aggregate
should be created using SyncMirror.
* block_storage.primary.checksum_style - Checksum style of the disks to be
use for the aggregate.
* block_storage.primary.disk_class - Class of disks to be use to for the
aggregate.
* block_storage.primary.raid_size - Desired RAID size of the aggregate.
* block_storage.primary.raid_type - Desired RAID type of the aggregate.
* snaplock_type - SnapLock type to use on the aggregate.
* data_encryption.software_encryption_enabled - Enable or disable NAE
(NetApp Aggregate Encryption) on the aggregate.
* simulate - Simulate the creation of the aggregate with specified input
parameters.
If the simulate field is specified, the response includes information on
the proposed aggregate disk layout, any associated warnings, along with
the proposed final size of the aggregate.

'''

== Examples

=== Retrieving a list of aggregates from the cluster

The following example shows the response with a list of data aggregates in
the cluster:

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates" -H "accept:
application/json"

# The response:
```

```
{
"records": [
  {
    "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
    "name": "test1",
    "node": {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
  },
  {
    "uuid": "4a7e4139-ca7a-420b-9a11-3f040d2189fd",
    "name": "test4",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2",
    },
  }
],
"num_records": 2,
}
----

=== Retrieving a list of aggregates recommended for creation from the
cluster

The following example shows the response with a list of recommended data
aggregates in the cluster.

NOTE: Each aggregate UUID provided in this response is not guaranteed to
be the same UUID for the aggregate if it is created.

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?recommend=true&fields=*" -H "accept:
application/json"

# The response:
{
"records": [
  {
```

```
    "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
    "name": "node_2_SSD_1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2",
    },
    "space": {
      "block_storage": {
        "size": 1116180480
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 22,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 24,
        "simulated_raid_groups": [
           {
             "name": "node_2_SSD_1/plex0/rg0",
             "raid_type": "raid_dp",
             "parity_disk_count": 2,
             "data_disk_count": 9,
             "usable_size": 12309487,
             "is_partition": true
           },
           {
             "name": "node_2_SSD_1/plex1/rg0",
             "raid_type": "raid_dp",
             "parity_disk_count": 2,
             "data_disk_count": 9,
             "usable_size": 12309487,
             "is_partition": false
           }
        ]
      },
      "hybrid_cache": {
        "enabled": true,
        "storage_pools": [
            {
                "allocation_units_count": 1,
                "storage_pool": {
                    "name": "sp1",
                    "uuid": "1511d084-7290-11ec-ae5b-005056bb2afa"
                }
```

```
            },
            {
                "allocation_units_count": 1,
                "storage_pool": {
                    "name": "sp2",
                    "uuid": "342d234f-7291-11ec-ae5b-005056bb2afa"
                }
            }
        ]
      },
      "mirror": {
        "enabled": true
      }
    },
  },
  {
    "uuid": "795c0a15-fa4b-11e8-ba65-005056bbe5c1",
    "name": "node_1_SSD_1",
    "node": {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
    "space": {
      "block_storage": {
        "size": 352477184
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 22,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 24,
        "simulated_raid_groups": [
           {
             "name": "node_1_SSD_1/plex0/rg0",
             "raid_type": "raid_dp",
             "parity_disk_count": 2,
             "data_disk_count": 9,
             "usable_size": 12309487,
             "is_partition": true
           },
           {
             "name": "node_1_SSD_1/plex1/rg0",
             "raid_type": "raid_dp",
```

```
              "parity_disk_count": 2,
              "data_disk_count": 9,
              "usable_size": 12309487,
              "is_partition": false
            }
          ]
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": true
      }
    },
  }
],
"recommendation_spares" [
  {
    "node": {
      "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5cl",
      "name": "node-2"
      }
    },
    "disk_class": "solid-state",
    "size": "2856845312",
    "is_partition": true,
    "disk_type": "ssd",
    "syncmirror_pool": "pool10",
    "usable": 1
  }
],
"num_records": 2,
"warnings": [
  {
    "name": "node_1_SSD_1",
    "warning": {
      "code": 19726347,
      "message": "Unable to use all attached capacity on node \"node_1\".
3 local/remote pool disks not usable for mirroring.",
      "arguments": [
        "node_1",
        "3"
      ]
    },
    "action": {
      "code": 19726348,
```

```
      "message": "Contact technical support."
    }
  },
  {
    "name": "node_2_SSD_1",
    "warning": {
      "code": 19726347,
      "message": "Unable to use all attached capacity on node \"node_2\".
3 local/remote pool disks not usable for mirroring.",
      "arguments": [
        "node_2",
        "3"
      ]
    },
    "action": {
      "code": 19726348,
      "message": "Contact technical support."
    }
  }
]
}
----

=== Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare
information according to ONTAP best practices.

----

 # The API:
/api/storage/aggregates?show_spares=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates?show_spares=true" -H
"accept: application/json"

# The response:
{
"records": [],
"num_records": 0,
"spares": [
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
```

```
    },
    "disk_class": "solid_state",
    "disk_type": "ssd",
    "size": 3720609792,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": true,
    "usable": 12,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  },
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
```

```
    "disk_class": "solid_state",
    "disk_type": "ssd_nvm",
    "size": 8583380992,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": false,
    "usable": 5,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  }
]
}
----

=== Retrieving the SSD spare count information for the cluster
```

```
The following example shows the response from retrieving SSD spare count
information.

----

# The API:
/api/storage/aggregates?show_spares=true&flash_pool_eligible=true

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?show_spares=true&flash_pool_eligible=true" -H
"accept: application/json"

# The response:
{
"records": [],
"num_records": 0,
"spares": [
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
    "disk_class": "solid_state",
    "disk_type": "ssd",
    "size": 3720609792,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": true,
    "usable": 12,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
```

```
          "min": 2,
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  },
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
    "disk_class": "solid_state",
    "disk_type": "ssd_nvm",
    "size": 8583380992,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": false,
    "usable": 5,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
```

```
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  }
]
}
----

=== Retrieving the total spare count information for the cluster

The following example shows the response from retrieving total spare count
information, under advanced privilege.

----

 # The API:
/api/storage/aggregates?show_spares=true&fields=**

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?show_spares=true&fields=**" -H "accept:
application/json"

# The response:
{
"records": [],
"num_records": 0,
"spares": [
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
    "disk_class": "solid_state",
```

```
    "disk_type": "ssd",
    "size": 3720609792,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": true,
    "usable": 12,
    "total": 14,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  },
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
    "disk_class": "solid_state",
```

```
    "disk_type": "ssd_nvm",
    "size": 8583380992,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": false,
    "usable": 5,
    "total": 6,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      },
      {
        "raid_type": "raid4",
        "default": false,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
          "max": 14,
          "default": 8
        }
      },
      {
        "raid_type": "raid_tec",
        "default": false,
        "aggregate_min_disks": 7,
        "raid_group": {
          "min": 4,
          "max": 29,
          "default": 25
        }
      }
    ]
  }
]
}
----

=== Simulating the creation of an aggregate using defined parameters
```

```
The following example shows the response in the simulation of a manual
aggregate creation:

----

# The API:
/api/storage/aggregates

# The call
curl -X POST "https://<mgmt-ip>/api/storage/aggregates" -H "accept:
application/json" -d "{\"name\":\"node_2_SSD_1\",\"node\":
{\"name\":\"node-2\"},\"block_storage\": {\"primary\": {\"disk_count\":
\"12\" }},\"simulate\": \"true\"}"

# The response:
{
"records": [
  {
    "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
    "name": "node_2_SSD_1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2",
    },
    "space": {
      "block_storage": {
        "size": 1116180480
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 12,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 24,
        "simulated_raid_groups": [
           {
             "name": "node_2_SSD_1/plex0/rg0",
             "raid_type": "raid_dp",
             "parity_disk_count": 2,
             "data_disk_count": 10,
             "usable_size": 558090240,
             "is_partition": true
           },
        ]
```

```
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": false
      }
    }
  }
],
"num_records": 1,
"warnings": [
  {
    "name": "node-2"
    "warning": {
      "code": 19726347,
      "message": "Number of unassigned disks attached to node \"node-2\":
6.",
      "arguments": [
        "node-2",
        "6"
      ]
    }
  }
]
}
```


```
----
```

```
[[ID9b632895a5a79c71982ffb6bfe5da915]]
= Retrieve a collection of aggregates for an entire cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates`#

*Introduced In:* 9.6

Retrieves the collection of aggregates for the entire cluster.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
```

```
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `metric.+*+`
* `space.block_storage.inactive_user_data`
* `space.block_storage.inactive_user_data_percent`
* `space.footprint`
* `is_spare_low`
* `statistics.+*+`

== Related ONTAP commands

* `storage aggregate show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|recommend
|boolean
|query
|False
a|If set to 'true', it queries the system for the recommended optimal
layout for creating new aggregates. The default setting is 'false'.


|show_spares
|boolean
|query
|False
a|If set to 'true', the spares object is returned instead of records to
show the spare groups in the cluster. The default setting is 'false'.


|flash_pool_eligible
|boolean
|query
|False
a|If set to 'true' along with show_spares, the spares object is restricted
```

```
to return spare groups that are compatible with whole disk based flash
pool creation. The default setting is 'false'.

* Introduced in: 9.12


|snapshot.files_total
|integer
|query
|False
a|Filter by snapshot.files_total

* Introduced in: 9.10


|snapshot.max_files_used
|integer
|query
|False
a|Filter by snapshot.max_files_used

* Introduced in: 9.10


|snapshot.files_used
|integer
|query
|False
a|Filter by snapshot.files_used

* Introduced in: 9.10


|snapshot.max_files_available
|integer
|query
|False
a|Filter by snapshot.max_files_available

* Introduced in: 9.10


|uuid
|string
|query
|False
a|Filter by uuid
```

```
|space.snapshot.total
|integer
|query
|False
a|Filter by space.snapshot.total

* Introduced in: 9.10


|space.snapshot.used_percent
|integer
|query
|False
a|Filter by space.snapshot.used_percent

* Introduced in: 9.10


|space.snapshot.reserve_percent
|integer
|query
|False
a|Filter by space.snapshot.reserve_percent

* Introduced in: 9.10


|space.snapshot.used
|integer
|query
|False
a|Filter by space.snapshot.used

* Introduced in: 9.10


|space.snapshot.available
|integer
|query
|False
a|Filter by space.snapshot.available

* Introduced in: 9.10
```

```
|space.efficiency_without_snapshots_flexclones.savings
|integer
|query
|False
a|Filter by space.efficiency_without_snapshots_flexclones.savings

* Introduced in: 9.9


|space.efficiency_without_snapshots_flexclones.logical_used
|integer
|query
|False
a|Filter by space.efficiency_without_snapshots_flexclones.logical_used

* Introduced in: 9.9


|space.efficiency_without_snapshots_flexclones.ratio
|number
|query
|False
a|Filter by space.efficiency_without_snapshots_flexclones.ratio

* Introduced in: 9.9


|space.total_provisioned_space
|integer
|query
|False
a|Filter by space.total_provisioned_space

* Introduced in: 9.18


|space.block_storage.physical_used
|integer
|query
|False
a|Filter by space.block_storage.physical_used

* Introduced in: 9.9


|space.block_storage.used_percent
|integer
```

```
|query
|False
a|Filter by space.block_storage.used_percent

* Introduced in: 9.13


|space.block_storage.volume_deduplication_shared_count
|integer
|query
|False
a|Filter by space.block_storage.volume_deduplication_shared_count

* Introduced in: 9.10


|space.block_storage.aggregate_metadata
|integer
|query
|False
a|Filter by space.block_storage.aggregate_metadata

* Introduced in: 9.10


|space.block_storage.used_including_snapshot_reserve_percent
|integer
|query
|False
a|Filter by space.block_storage.used_including_snapshot_reserve_percent

* Introduced in: 9.10


|space.block_storage.volume_footprints_percent
|integer
|query
|False
a|Filter by space.block_storage.volume_footprints_percent

* Introduced in: 9.10


|space.block_storage.inactive_user_data
|integer
|query
|False
```

```
a|Filter by space.block_storage.inactive_user_data


|space.block_storage.used
|integer
|query
|False
a|Filter by space.block_storage.used


|space.block_storage.volume_deduplication_space_saved
|integer
|query
|False
a|Filter by space.block_storage.volume_deduplication_space_saved

* Introduced in: 9.10


|space.block_storage.data_compacted_count
|integer
|query
|False
a|Filter by space.block_storage.data_compacted_count

* Introduced in: 9.10


|space.block_storage.size
|integer
|query
|False
a|Filter by space.block_storage.size


|space.block_storage.available
|integer
|query
|False
a|Filter by space.block_storage.available


|space.block_storage.aggregate_metadata_percent
|integer
|query
|False
a|Filter by space.block_storage.aggregate_metadata_percent
```

```
* Introduced in: 9.10


|space.block_storage.inactive_user_data_percent
|integer
|query
|False
a|Filter by space.block_storage.inactive_user_data_percent

* Introduced in: 9.10


|space.block_storage.data_compaction_space_saved
|integer
|query
|False
a|Filter by space.block_storage.data_compaction_space_saved

* Introduced in: 9.10


|space.block_storage.data_compaction_space_saved_percent
|integer
|query
|False
a|Filter by space.block_storage.data_compaction_space_saved_percent

* Introduced in: 9.10


|space.block_storage.physical_used_percent
|integer
|query
|False
a|Filter by space.block_storage.physical_used_percent

* Introduced in: 9.10


|space.block_storage.used_including_snapshot_reserve
|integer
|query
|False
a|Filter by space.block_storage.used_including_snapshot_reserve

* Introduced in: 9.10
```

```
|space.block_storage.volume_deduplication_space_saved_percent
|integer
|query
|False
a|Filter by space.block_storage.volume_deduplication_space_saved_percent

* Introduced in: 9.10


|space.block_storage.full_threshold_percent
|integer
|query
|False
a|Filter by space.block_storage.full_threshold_percent


|space.efficiency.cross_volume_inline_dedupe
|boolean
|query
|False
a|Filter by space.efficiency.cross_volume_inline_dedupe

* Introduced in: 9.12


|space.efficiency.ratio
|number
|query
|False
a|Filter by space.efficiency.ratio


|space.efficiency.enable_workload_informed_tsse
|boolean
|query
|False
a|Filter by space.efficiency.enable_workload_informed_tsse

* Introduced in: 9.13


|space.efficiency.wise_tsse_min_used_capacity_pct
|integer
|query
|False
```

```
a|Filter by space.efficiency.wise_tsse_min_used_capacity_pct

* Introduced in: 9.13


|space.efficiency.auto_adaptive_compression_savings
|boolean
|query
|False
a|Filter by space.efficiency.auto_adaptive_compression_savings

* Introduced in: 9.12


|space.efficiency.savings
|integer
|query
|False
a|Filter by space.efficiency.savings


|space.efficiency.cross_volume_background_dedupe
|boolean
|query
|False
a|Filter by space.efficiency.cross_volume_background_dedupe

* Introduced in: 9.12


|space.efficiency.cross_volume_dedupe_savings
|boolean
|query
|False
a|Filter by space.efficiency.cross_volume_dedupe_savings

* Introduced in: 9.12


|space.efficiency.logical_used
|integer
|query
|False
a|Filter by space.efficiency.logical_used


|space.efficiency_without_snapshots.savings
```

```
|integer
|query
|False
a|Filter by space.efficiency_without_snapshots.savings


|space.efficiency_without_snapshots.logical_used
|integer
|query
|False
a|Filter by space.efficiency_without_snapshots.logical_used


|space.efficiency_without_snapshots.ratio
|number
|query
|False
a|Filter by space.efficiency_without_snapshots.ratio


|space.cloud_storage.used
|integer
|query
|False
a|Filter by space.cloud_storage.used


|space.footprint
|integer
|query
|False
a|Filter by space.footprint


|home_node.name
|string
|query
|False
a|Filter by home_node.name


|home_node.uuid
|string
|query
|False
a|Filter by home_node.uuid
```

```
|inactive_data_reporting.start_time
|string
|query
|False
a|Filter by inactive_data_reporting.start_time

* Introduced in: 9.8


|inactive_data_reporting.enabled
|boolean
|query
|False
a|Filter by inactive_data_reporting.enabled

* Introduced in: 9.8


|metric.status
|string
|query
|False
a|Filter by metric.status

* Introduced in: 9.7


|metric.duration
|string
|query
|False
a|Filter by metric.duration

* Introduced in: 9.7


|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other

* Introduced in: 9.7


|metric.latency.total
```

```
|integer
|query
|False
a|Filter by metric.latency.total

* Introduced in: 9.7


|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write

* Introduced in: 9.7


|metric.latency.read
|integer
|query
|False
a|Filter by metric.latency.read

* Introduced in: 9.7


|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other

* Introduced in: 9.7


|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total

* Introduced in: 9.7


|metric.iops.write
|integer
|query
```

```
|False
a|Filter by metric.iops.write

* Introduced in: 9.7


|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read

* Introduced in: 9.7


|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp

* Introduced in: 9.7


|metric.throughput.other
|integer
|query
|False
a|Filter by metric.throughput.other

* Introduced in: 9.7


|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total

* Introduced in: 9.7


|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

```
* Introduced in: 9.7


|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read

* Introduced in: 9.7


|inode_attributes.used_percent
|integer
|query
|False
a|Filter by inode_attributes.used_percent

* Introduced in: 9.11
* Max value: 100
* Min value: 0


|inode_attributes.file_private_capacity
|integer
|query
|False
a|Filter by inode_attributes.file_private_capacity

* Introduced in: 9.11


|inode_attributes.version
|integer
|query
|False
a|Filter by inode_attributes.version

* Introduced in: 9.11


|inode_attributes.files_used
|integer
|query
|False
a|Filter by inode_attributes.files_used
```

```
* Introduced in: 9.11


|inode_attributes.files_private_used
|integer
|query
|False
a|Filter by inode_attributes.files_private_used

* Introduced in: 9.11


|inode_attributes.max_files_available
|integer
|query
|False
a|Filter by inode_attributes.max_files_available

* Introduced in: 9.11


|inode_attributes.file_public_capacity
|integer
|query
|False
a|Filter by inode_attributes.file_public_capacity

* Introduced in: 9.11


|inode_attributes.files_total
|integer
|query
|False
a|Filter by inode_attributes.files_total

* Introduced in: 9.11


|inode_attributes.max_files_used
|integer
|query
|False
a|Filter by inode_attributes.max_files_used

* Introduced in: 9.11
```

```
|inode_attributes.max_files_possible
|integer
|query
|False
a|Filter by inode_attributes.max_files_possible

* Introduced in: 9.11


|cloud_storage.attach_eligible
|boolean
|query
|False
a|Filter by cloud_storage.attach_eligible

* Introduced in: 9.14


|cloud_storage.stores.used
|integer
|query
|False
a|Filter by cloud_storage.stores.used

* Introduced in: 9.14


|cloud_storage.stores.cloud_store.uuid
|string
|query
|False
a|Filter by cloud_storage.stores.cloud_store.uuid

* Introduced in: 9.14


|cloud_storage.stores.cloud_store.name
|string
|query
|False
a|Filter by cloud_storage.stores.cloud_store.name

* Introduced in: 9.14
```

```
|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|dr_home_node.name
|string
|query
|False
a|Filter by dr_home_node.name


|dr_home_node.uuid
|string
|query
|False
a|Filter by dr_home_node.uuid


|snaplock_type
|string
|query
|False
a|Filter by snaplock_type


|sidl_enabled
|boolean
|query
|False
a|Filter by sidl_enabled

* Introduced in: 9.11


|statistics.status
|string
|query
```

```
|False
a|Filter by statistics.status

* Introduced in: 9.7


|statistics.throughput_raw.other
|integer
|query
|False
a|Filter by statistics.throughput_raw.other

* Introduced in: 9.7


|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total

* Introduced in: 9.7


|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.7


|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read

* Introduced in: 9.7


|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other
```

```
* Introduced in: 9.7


|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total

* Introduced in: 9.7


|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write

* Introduced in: 9.7


|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read

* Introduced in: 9.7


|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

* Introduced in: 9.7


|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total

* Introduced in: 9.7
```

```
|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

* Introduced in: 9.7


|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read

* Introduced in: 9.7


|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp

* Introduced in: 9.7


|volume-count
|integer
|query
|False
a|Filter by volume-count

* Introduced in: 9.11


|data_encryption.software_encryption_enabled
|boolean
|query
|False
a|Filter by data_encryption.software_encryption_enabled


|data_encryption.drive_protection_enabled
|boolean
```

```
|query
|False
a|Filter by data_encryption.drive_protection_enabled


|state
|string
|query
|False
a|Filter by state


|name
|string
|query
|False
a|Filter by name


|create_time
|string
|query
|False
a|Filter by create_time


|block_storage.hybrid_cache.storage_pools.allocation_units_count
|integer
|query
|False
a|Filter by
block_storage.hybrid_cache.storage_pools.allocation_units_count

* Introduced in: 9.11


|block_storage.hybrid_cache.storage_pools.storage_pool.name
|string
|query
|False
a|Filter by block_storage.hybrid_cache.storage_pools.storage_pool.name

* Introduced in: 9.11


|block_storage.hybrid_cache.storage_pools.storage_pool.uuid
|string
```

```
|query
|False
a|Filter by block_storage.hybrid_cache.storage_pools.storage_pool.uuid

* Introduced in: 9.11


|block_storage.hybrid_cache.simulated_raid_groups.usable_size
|integer
|query
|False
a|Filter by block_storage.hybrid_cache.simulated_raid_groups.usable_size

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.added_parity_disk_count
|integer
|query
|False
a|Filter by
block_storage.hybrid_cache.simulated_raid_groups.added_parity_disk_count

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.name
|string
|query
|False
a|Filter by block_storage.hybrid_cache.simulated_raid_groups.name

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.added_data_disk_count
|integer
|query
|False
a|Filter by
block_storage.hybrid_cache.simulated_raid_groups.added_data_disk_count

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.is_partition
|boolean
```

```
|query
|False
a|Filter by block_storage.hybrid_cache.simulated_raid_groups.is_partition

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.existing_data_disk_count
|integer
|query
|False
a|Filter by
block_storage.hybrid_cache.simulated_raid_groups.existing_data_disk_count

* Introduced in: 9.12


|block_storage.hybrid_cache.simulated_raid_groups.existing_parity_disk_cou
nt
|integer
|query
|False
a|Filter by
block_storage.hybrid_cache.simulated_raid_groups.existing_parity_disk_coun
t

* Introduced in: 9.12


|block_storage.hybrid_cache.enabled
|boolean
|query
|False
a|Filter by block_storage.hybrid_cache.enabled


|block_storage.hybrid_cache.raid_size
|integer
|query
|False
a|Filter by block_storage.hybrid_cache.raid_size

* Introduced in: 9.12


|block_storage.hybrid_cache.disk_count
|integer
```

```
|query
|False
a|Filter by block_storage.hybrid_cache.disk_count


|block_storage.hybrid_cache.used
|integer
|query
|False
a|Filter by block_storage.hybrid_cache.used


|block_storage.hybrid_cache.disk_type
|string
|query
|False
a|Filter by block_storage.hybrid_cache.disk_type

* Introduced in: 9.12


|block_storage.hybrid_cache.size
|integer
|query
|False
a|Filter by block_storage.hybrid_cache.size


|block_storage.hybrid_cache.raid_type
|string
|query
|False
a|Filter by block_storage.hybrid_cache.raid_type


|block_storage.storage_type
|string
|query
|False
a|Filter by block_storage.storage_type

* Introduced in: 9.11


|block_storage.mirror.state
|string
|query
```

```
|False
a|Filter by block_storage.mirror.state


|block_storage.mirror.enabled
|boolean
|query
|False
a|Filter by block_storage.mirror.enabled


|block_storage.primary.checksum_style
|string
|query
|False
a|Filter by block_storage.primary.checksum_style


|block_storage.primary.disk_class
|string
|query
|False
a|Filter by block_storage.primary.disk_class


|block_storage.primary.disk_count
|integer
|query
|False
a|Filter by block_storage.primary.disk_count


|block_storage.primary.raid_type
|string
|query
|False
a|Filter by block_storage.primary.raid_type


|block_storage.primary.raid_size
|integer
|query
|False
a|Filter by block_storage.primary.raid_size


|block_storage.primary.simulated_raid_groups.name
```

```
|string
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.name

* Introduced in: 9.10


|block_storage.primary.simulated_raid_groups.existing_parity_disk_count
|integer
|query
|False
a|Filter by
block_storage.primary.simulated_raid_groups.existing_parity_disk_count

* Introduced in: 9.11


|block_storage.primary.simulated_raid_groups.added_parity_disk_count
|integer
|query
|False
a|Filter by
block_storage.primary.simulated_raid_groups.added_parity_disk_count

* Introduced in: 9.11


|block_storage.primary.simulated_raid_groups.parity_disk_count
|integer
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.parity_disk_count

* Introduced in: 9.10


|block_storage.primary.simulated_raid_groups.added_data_disk_count
|integer
|query
|False
a|Filter by
block_storage.primary.simulated_raid_groups.added_data_disk_count

* Introduced in: 9.11
```

```
|block_storage.primary.simulated_raid_groups.is_partition
|boolean
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.is_partition

* Introduced in: 9.10


|block_storage.primary.simulated_raid_groups.existing_data_disk_count
|integer
|query
|False
a|Filter by
block_storage.primary.simulated_raid_groups.existing_data_disk_count

* Introduced in: 9.11


|block_storage.primary.simulated_raid_groups.raid_type
|string
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.raid_type

* Introduced in: 9.10


|block_storage.primary.simulated_raid_groups.usable_size
|integer
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.usable_size

* Introduced in: 9.10


|block_storage.primary.simulated_raid_groups.data_disk_count
|integer
|query
|False
a|Filter by block_storage.primary.simulated_raid_groups.data_disk_count

* Introduced in: 9.10


|block_storage.primary.disk_type
```

```
|string
|query
|False
a|Filter by block_storage.primary.disk_type

* Introduced in: 9.7


|block_storage.uses_partitions
|boolean
|query
|False
a|Filter by block_storage.uses_partitions

* Introduced in: 9.11


|block_storage.plexes.name
|string
|query
|False
a|Filter by block_storage.plexes.name


|_tags
|string
|query
|False
a|Filter by _tags

* Introduced in: 9.13


|is_spare_low
|boolean
|query
|False
a|Filter by is_spare_low

* Introduced in: 9.11


|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records


|recommendation_spares
|array[link:#aggregate_spare[aggregate_spare]]
a|Information on the aggregate's remaining hot spare disks.


|records
|array[link:#aggregate[aggregate]]
a|

|spares
|array[link:#aggregate_spare[aggregate_spare]]
a|

|warnings
|array[link:#aggregate_warning[aggregate_warning]]
a|List of warnings and remediation advice for the aggregate
recommendation.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
```

```
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "recommendation_spares": [
    {
      "checksum_style": "string",
      "disk_class": "solid_state",
      "disk_type": "string",
      "is_partition": 1,
      "layout_requirements": [
        {
          "aggregate_min_disks": 6,
          "raid_group": {
            "default": 16,
            "max": 28,
            "min": 5
          },
          "raid_type": "string"
        }
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 10156769280,
      "syncmirror_pool": "string",
```

```
      "total": 10,
      "usable": 9
    }
  ],
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "block_storage": {
        "hybrid_cache": {
          "disk_count": 6,
          "disk_type": "string",
          "raid_size": 24,
          "raid_type": "string",
          "simulated_raid_groups": [
            {
              "name": "string"
            }
          ],
          "size": 1612709888,
          "storage_pools": [
            {
              "storage_pool": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "storage_pool_1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            }
          ],
          "used": 26501122
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
```

```
      "plexes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "plex0"
        }
      ],
      "primary": {
        "checksum_style": "string",
        "disk_class": "performance",
        "disk_count": 8,
        "disk_type": "string",
        "raid_size": 16,
        "raid_type": "string",
        "simulated_raid_groups": [
          {
            "name": "string",
            "raid_type": "string"
          }
        ]
      },
      "storage_type": "string"
    },
    "cloud_storage": {
      "stores": [
        {
          "cloud_store": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "store1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "used": 0
        }
      ]
    },
    "create_time": "2018-01-01 12:00:00 -0400",
    "dr_home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "inactive_data_reporting": {
        "start_time": "2019-12-12 12:00:00 -0400"
      },
      "inode_attributes": {
        "file_private_capacity": 31136,
        "file_public_capacity": 31136,
        "files_private_used": 502,
        "files_total": 31136,
        "files_used": 97,
        "max_files_available": 31136,
        "max_files_possible": 2844525,
        "max_files_used": 97,
        "used_percent": 5,
        "version": 4
      },
      "is_spare_low": "",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
```

```
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "node1_aggr_1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snaplock_type": "string",
      "snapshot": {
        "files_total": 10,
        "files_used": 3,
        "max_files_available": 5,
        "max_files_used": 50
      },
      "space": {
        "block_storage": {
          "aggregate_metadata": 2655,
          "aggregate_metadata_percent": 8,
          "available": 10156560384,
          "data_compacted_count": 1990000,
          "data_compaction_space_saved": 1996000,
          "data_compaction_space_saved_percent": 27,
          "full_threshold_percent": 0,
          "inactive_user_data": 304448,
          "inactive_user_data_percent": 0,
          "physical_used": 2461696,
          "physical_used_percent": 50,
          "size": 10156769280,
          "used": 2088960,
          "used_including_snapshot_reserve": 674685,
          "used_including_snapshot_reserve_percent": 35,
          "used_percent": 50,
          "volume_deduplication_shared_count": 1990000,
          "volume_deduplication_space_saved": 1996000,
          "volume_deduplication_space_saved_percent": 27,
          "volume_footprints_percent": 14
        },
        "cloud_storage": {
```

```
          "used": 402743264
        },
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0,
          "wise_tsse_min_used_capacity_pct": 0
        },
        "efficiency_without_snapshots": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots_flexclones": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "footprint": 608896,
        "snapshot": {
          "available": 2000,
          "reserve_percent": 20,
          "total": 5000,
          "used": 3000,
          "used_percent": 45
        },
        "total_provisioned_space": 20971520
      },
      "state": "string",
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
```

```
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "uuid": "string",
      "volume-count": 0
    }
  ],
  "spares": [
    {
      "checksum_style": "string",
      "disk_class": "solid_state",
      "disk_type": "string",
      "is_partition": 1,
      "layout_requirements": [
        {
          "aggregate_min_disks": 6,
          "raid_group": {
            "default": 16,
            "max": 28,
            "min": 5
          },
          "raid_type": "string"
        }
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 10156769280,
      "syncmirror_pool": "string",
      "total": 10,
      "usable": 9
    }
  ],
  "warnings": [
    {
      "action": {
        "arguments": [
          "string"
        ],
        "message": "string"
      },
```

```
      "name": "string",
      "warning": {
        "arguments": [
          "string"
        ],
        "message": "string"
      }
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 787092
| The target field cannot be specified for this operation.

| 918138
| Internal error. Failed to get encryption operation status.

| 8586225
| Encountered unexpected error in retrieving metrics and statistics for an
aggregate.

| 19726341
| Not enough eligible spare disks are available on the node.

| 19726344
| No recommendation can be made for this cluster.

| 19726357
| Aggregate recommendations are not supported on MetroCluster with Fibre
Channel (FC).

| 19726358
| Aggregate recommendations are not supported on ONTAP Cloud.

| 19726382
| Another provisioning operation is in progress on this cluster. Wait a
```

```
few minutes, and try the operation again.

| 19726386
| Encountered an error when retrieving licensing information on this
cluster.

| 19726387
| No recommendation can be provided for this cluster within the license
capacity.

| 19726401
| Aggregate recommendations are not supported when the DR group is not in
the "normal" state.

| 19726402
| Internal error. Unable to determine the MetroCluster configuration
state.

| 19726403
| Aggregate recommendation is not supported when there are no healthy
target connections to remote storage.

| 19726404
| The recommended mirrored aggregate couldn't use all the attached
capacity in one of the SyncMirror pools. Make sure that the remote and
local storage is symmetrically wired.

| 19726405
| Not all local and remote disks attached to the node have been auto-
partitioned.

| 19726406
| Aggregate recommendations are not supported on this node because remote
and local storage is not symmetrically wired.

| 19726540
| The next tag is not supported for recommended aggregates. Retry the
operation with a higher "return_timeout" value.

| 196608055
| Aggregate recommendation is not supported on this node because it does
not support NetApp Aggregate Encryption (NAE).

| 196608206
| Internal error. Failed to get encryption operation status.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#raid_group]
[.api-collapsible-fifth-title]
raid_group

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|default
|integer
a|Default number of disks in a RAID group.


|max
|integer
a|Maximum number of disks allowed in a RAID group.


|min
|integer
a|Minimum number of disks allowed in a RAID group.


|===


[#layout_requirement]
[.api-collapsible-fifth-title]
layout_requirement

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate_min_disks
|integer
a|Minimum number of disks to create an aggregate.


|default
|boolean
a|Indicates if this RAID type is the default.


|raid_group
|link:#raid_group[raid_group]
a|

|raid_type
|string
```

```
a|RAID type.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

Node where the spares are assigned.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#aggregate_spare]
[.api-collapsible-fifth-title]
aggregate_spare

[cols=3*,options=header]
|===
|Name
|Type
|Description

|checksum_style
|string
a|The checksum type that has been assigned to the spares.


|disk_class
|string
a|Disk class of spares.


|disk_type
|string
a|Type of disk.


|is_partition
|boolean
a|Indicates whether a disk is partitioned (true) or whole (false)


|layout_requirements
|array[link:#layout_requirement[layout_requirement]]
a|Available RAID protections and their restrictions.


|node
|link:#node[node]
a|Node where the spares are assigned.


|size
|integer
a|Usable size of each spare, in bytes.
```

```
|syncmirror_pool
|string
a|SyncMirror spare pool.


|total
|integer
a|Total number of spares in the bucket. The total spare count for each
class of spares also includes reserved spare capacity recommended by ONTAP
best practices.

* example: 10
* readOnly: 1
* Introduced in: 9.11
* x-nullable: true


|usable
|integer
a|Total number of usable spares in the bucket. The usable count for each
class of spares does not include reserved spare capacity recommended by
ONTAP best practices.

* example: 9
* readOnly: 1
* Introduced in: 9.6
* x-nullable: true


|===


[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.
```

```
|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#storage_pool_reference]
[.api-collapsible-fifth-title]
storage_pool_reference

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pools]
[.api-collapsible-fifth-title]
storage_pools

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allocation_units_count
|integer
a|Allocation count of storage pool.


|storage_pool
|link:#storage_pool_reference[storage_pool_reference]
a|Shared Storage Pool


|===


[#hybrid_cache]
[.api-collapsible-fifth-title]
hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made
up of either whole SSDs or storage pool SSDs.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|Number of disks used in the cache tier of the aggregate. Only provided
when hybrid_cache.enabled is 'true'.


|disk_type
|string
a|Type of disk being used by the aggregate's cache tier.


|enabled
|boolean
a|Specifies whether the aggregate uses HDDs with SSDs as a cache.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type for SSD cache of the aggregate. Only provided when
hybrid_cache.enabled is 'true'.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|size
|integer
a|Total usable space in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|storage_pools
|array[link:#storage_pools[storage_pools]]
a|List of storage pool properties and allocation_units_count for
```

```
aggregate.


|used
|integer
a|Space used in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|===


[#mirror]
[.api-collapsible-fifth-title]
mirror

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Aggregate is SyncMirror protected


|state
|string
a|

|===


[#plex_reference]
[.api-collapsible-fifth-title]
plex_reference

Plex


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|===


[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|data_disk_count
|integer
a|Number of data disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
```

```
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|parity_disk_count
|integer
a|Number of parity disks in RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#primary]
[.api-collapsible-fifth-title]
primary

Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|checksum_style
|string
a|The checksum style used by the aggregate.


|disk_class
```

```
|string
a|The class of disks being used by the aggregate.


|disk_count
|integer
a|Number of disks used in the aggregate. This includes parity disks, but
excludes disks in the hybrid cache.


|disk_type
|string
a|The type of disk being used by the aggregate.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|hybrid_cache
|link:#hybrid_cache[hybrid_cache]
a|Contains the configuration for the hybrid cache. The hybrid cache is
made up of either whole SSDs or storage pool SSDs.


|mirror
|link:#mirror[mirror]
a|

|plexes
|array[link:#plex_reference[plex_reference]]
a|Plex reference for each plex in the aggregate.


|primary
|link:#primary[primary]
a|Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


|storage_type
|string
a|Type of aggregate.


|uses_partitions
|boolean
a|If true, aggregate is using shared disks.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

Cloud store


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#cloud_storage_tier]
[.api-collapsible-fifth-title]
cloud_storage_tier

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cloud_store
|link:#cloud_store[cloud_store]
a|Cloud store


|used
|integer
a|Capacity used in bytes in the cloud store by this aggregate. This is a
cached value calculated every 5 minutes.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

Configuration information for the cloud storage portion of the aggregate.
This is referred to as the capacity tier.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|attach_eligible
|boolean
a|Specifies whether the aggregate is eligible for a cloud store to be
attached.


|migrate_threshold
|integer
a|Specifies the minimum percentage of performance tier free space that
must exist in order for migration of data from the capacity tier to
performance tier to be allowed. Only valid for PATCH operations.


|stores
|array[link:#cloud_storage_tier[cloud_storage_tier]]
a|Configuration information for each cloud storage portion of the
aggregate.


|tiering_fullness_threshold
|integer
a|The percentage of space in the performance tier that must be used before
data is tiered out to the cloud store. Only valid for PATCH operations.


|===


[#data_encryption]
[.api-collapsible-fifth-title]
data_encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drive_protection_enabled
|boolean
```

```
a|Specifies whether the aggregate uses self-encrypting drives with data
protection enabled.


|software_encryption_enabled
|boolean
a|Specifies whether NetApp aggregate encryption is enabled. All data in
the aggregate is encrypted.


|===


[#dr_home_node]
[.api-collapsible-fifth-title]
dr_home_node

Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

Node where the aggregate resides after giveback. The value for this field
might differ from the value of the 'node' field during takeover.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#inactive_data_reporting]
[.api-collapsible-fifth-title]
inactive_data_reporting

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether or not inactive data reporting is enabled on the
aggregate.


|start_time
|string
a|Timestamp at which inactive data reporting was enabled on the aggregate.


|===


[#inode_attributes]
[.api-collapsible-fifth-title]
inode_attributes
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|file_private_capacity
|integer
a|Number of files that can currently be stored on disk for system metadata
files. This number will dynamically increase as more system files are
created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|file_public_capacity
|integer
a|Number of files that can currently be stored on disk for user-visible
files.  This number will dynamically increase as more user-visible files
are created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_private_used
|integer
a|Number of system metadata files used. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_total
|integer
a|Maximum number of user-visible files that this referenced file system
can currently hold. If the referenced file system is restricted or
offline, a value of 0 is returned.
```

```
|files_used
|integer
a|Number of user-visible files used in the referenced file system. If the
referenced file system is restricted or offline, a value of 0 is returned.


|max_files_available
|integer
a|The count of the maximum number of user-visible files currently
allowable on the referenced file system.


|max_files_possible
|integer
a|The largest value to which the maxfiles-available parameter can be
increased by reconfiguration, on the referenced file system.


|max_files_used
|integer
a|The number of user-visible files currently in use on the referenced file
system.


|used_percent
|integer
a|The percentage of disk space currently in use based on user-visible file
count on the referenced file system.


|version
|integer
a|The inofile-version of the aggregate. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|===


[#iops]
[.api-collapsible-fifth-title]
iops
```

```
The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

The most recent sample of I/O metrics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node where the aggregate currently resides.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

```
a|

|uuid
|string
a|

|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|files_total
|integer
a|Total files allowed in snapshots


|files_used
|integer
a|Total files created in snapshots


|max_files_available
|integer
a|Maximum files available for snapshots


|max_files_used
|integer
a|Files in use by snapshots


|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|aggregate_metadata
|integer
a|Space used by different metafiles and internal operations inside the
aggregate, in bytes.


|aggregate_metadata_percent
|integer
a|Aggregate metadata as a percentage.


|available
|integer
a|Space available in bytes.


|data_compacted_count
|integer
a|Amount of compacted data in bytes.


|data_compaction_space_saved
|integer
a|Space saved in bytes by compacting the data.


|data_compaction_space_saved_percent
|integer
a|Percentage saved by compacting the data.


|full_threshold_percent
|integer
a|The aggregate used percentage at which 'monitor.volume.full' EMS is
generated.


|inactive_user_data
|integer
a|The size that is physically used in the block storage and has a cold
temperature, in bytes. This property is only supported if the aggregate is
either attached to a cloud store or can be attached to a cloud store.
```

```
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data or **.


|inactive_user_data_percent
|integer
a|The percentage of inactive user data in the block storage. This property
is only supported if the aggregate is either attached to a cloud store or
can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data_percent or **.


|physical_used
|integer
a|Total physical used size of an aggregate in bytes.


|physical_used_percent
|integer
a|Physical used percentage.


|size
|integer
a|Total usable space in bytes, not including WAFL reserve and aggregate
snapshot reserve.


|used
|integer
a|Space used or reserved in bytes. Includes volume guarantees and
aggregate metadata.


|used_including_snapshot_reserve
|integer
a|Total used including the snapshot reserve, in bytes.


|used_including_snapshot_reserve_percent
```

```
|integer
a|Total used including the Snapshot reserve as a percentage.


|used_percent
|integer
a|Aggregate used percentage.


|volume_deduplication_shared_count
|integer
a|Amount of shared bytes counted by storage efficiency.


|volume_deduplication_space_saved
|integer
a|Amount of space saved in bytes by storage efficiency.


|volume_deduplication_space_saved_percent
|integer
a|Percentage of space saved by storage efficiency.


|volume_footprints_percent
|integer
a|A summation of volume footprints inside the aggregate, as a percentage.
A volume's footprint is the amount of space being used for the volume in
the aggregate.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|used
|integer
a|Used space in bytes in the cloud store. Only applicable for aggregates
```

```
with a cloud store tier.


|===


[#efficiency]
[.api-collapsible-fifth-title]
efficiency

Storage efficiency.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_adaptive_compression_savings
|boolean
a|Indicates whether or not aggregate has auto adaptive compression
savings.


|cross_volume_background_dedupe
|boolean
a|Indicates whether or not cross volume background deduplication is
enabled.


|cross_volume_dedupe_savings
|boolean
a|Indicates whether or not aggregate has cross volume deduplication
savings.


|cross_volume_inline_dedupe
|boolean
a|Indicates whether or not cross volume inline deduplication is enabled.


|enable_workload_informed_tsse
|boolean
a|Indicates whether Workload Informed TSSE is enabled on the system.
```

```
|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|wise_tsse_min_used_capacity_pct
|integer
a|Minimum amount of used data in aggregate required to trigger cold
compression on TSSE volumes.


|===


[#efficiency_without_snapshots]
[.api-collapsible-fifth-title]
efficiency_without_snapshots

Storage efficiency that does not include the savings provided by
snapshots.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)
```

```
|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots_flexclones]
[.api-collapsible-fifth-title]
efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by snapshots
and Flexclone volumes.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|available
|integer
a|Available space for snapshots in bytes


|reserve_percent
|integer
a|Percentage of space reserved for snapshots


|total
|integer
a|Total space for snapshots in bytes


|used
|integer
a|Space used by snapshots in bytes


|used_percent
|integer
a|Percentage of disk space used by snapshots


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_storage
|link:#block_storage[block_storage]
a|

|cloud_storage
```

```
|link:#cloud_storage[cloud_storage]
a|

|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and Flexclone volumes.


|footprint
|integer
a|A summation of volume footprints (including volume guarantees), in
bytes. This includes all of the volume footprints in the block_storage
tier and the cloud_storage tier.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|snapshot
|link:#snapshot[snapshot]
a|

|total_provisioned_space
|integer
a|A summation of volume size, in bytes. This includes all of the volumes
inside the aggregate.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either total_provisioned_space or **.


|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
```

```
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

The real time I/O statistics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
```

```
a|The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|cloud_storage
|link:#cloud_storage[cloud_storage]
a|Configuration information for the cloud storage portion of the
aggregate. This is referred to as the capacity tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|

|dr_home_node
|link:#dr_home_node[dr_home_node]
a|Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.


|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|

|inode_attributes
```

```
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the aggregate.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snaplock_type
|string
a|SnapLock type.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
```

```
|string
a|Operational state of the aggregate.


|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the aggregate.


|uuid
|string
a|Aggregate UUID.


|volume-count
|integer
a|Number of volumes in the aggregate.


|===


[#action]
[.api-collapsible-fifth-title]
action

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the specified action message.


|code
|integer
a|Corrective action code of the specified action.


|message
|string
a|Specifies the corrective action to be taken to resolve the issue.
```

```
|===


[#warning]
[.api-collapsible-fifth-title]
warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the warning message encountered.


|code
|integer
a|Warning code of the warning encountered.


|message
|string
a|Details of the warning encountered by the aggregate simulate query.


|===


[#aggregate_warning]
[.api-collapsible-fifth-title]
aggregate_warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|link:#action[action]
a|

|name
|string
```

```
a|Name of the entity that returns the warning.


|warning
|link:#warning[warning]
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID7abeeef124376447cf4185ae1fe0f298]]
= Create a collection of aggregates for an entire cluster

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/aggregates`#

*Introduced In:* 9.6

Automatically creates aggregates based on an optimal layout recommended by
the system. Alternatively, properties can be provided to create an
aggregate according to the requested specification. This request starts a
job and returns a link to that job.
POST operations will be blocked while one or more nodes in the cluster are
simulating or implementing automatic aggregate creation.

== Required properties

Properties are not required for this API. The following properties are
only required if you want to specify properties for aggregate creation:

* `name` - Name of the aggregate.
* `node.name` or `node.uuid` - Node on which the aggregate will be
created.
* `block_storage.primary.disk_count` - Number of disks to be used to
create the aggregate.

== Default values

If not specified in POST, the following default values are assigned. The
remaining unspecified properties will receive system dependent default
values.

* `block_storage.mirror.enabled` - _false_
* `snaplock_type` - _non_snaplock_

== Related ONTAP commands

* `storage aggregate auto-provision`
* `storage aggregate create`

== Example:

----
POST /api/storage/aggregates {"node": {"name": "node1"}, "name": "test",
"block_storage": {"primary": {"disk_count": "10"}}}
----
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|disk_size
|integer
|query
|False
a|If set, POST only selects disks of the specified size.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|

|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|

|inode_attributes
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
```

```
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snaplock_type
|string
a|SnapLock type.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
|string
a|Operational state of the aggregate.


|uuid
|string
a|Aggregate UUID.
```

```
|volume-count
|integer
a|Number of volumes in the aggregate.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "string",
      "raid_size": 24,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string"
        }
      ],
      "size": 1612709888,
      "storage_pools": [
        {
          "storage_pool": {
            "name": "storage_pool_1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "name": "plex0"
```

```
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "disk_type": "string",
      "raid_size": 16,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string",
          "raid_type": "string"
        }
      ]
    },
    "storage_type": "string"
  },
  "create_time": "2018-01-01 12:00:00 -0400",
  "home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "inactive_data_reporting": {
    "enabled": true,
    "start_time": "2019-12-12 12:00:00 -0400"
  },
  "inode_attributes": {
    "file_private_capacity": 31136,
    "file_public_capacity": 31136,
    "files_private_used": 502,
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "name": "node1_aggr_1",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snaplock_type": "string",
```

```
  "snapshot": {
    "files_total": 10,
    "files_used": 3,
    "max_files_available": 5,
    "max_files_used": 50
  },
  "space": {
    "block_storage": {
      "aggregate_metadata": 2655,
      "aggregate_metadata_percent": 8,
      "available": 10156560384,
      "data_compacted_count": 1990000,
      "data_compaction_space_saved": 1996000,
      "data_compaction_space_saved_percent": 27,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "inactive_user_data_percent": 0,
      "physical_used": 2461696,
      "physical_used_percent": 50,
      "size": 10156769280,
      "used": 2088960,
      "used_including_snapshot_reserve": 674685,
      "used_including_snapshot_reserve_percent": 35,
      "used_percent": 50,
      "volume_deduplication_shared_count": 1990000,
      "volume_deduplication_space_saved": 1996000,
      "volume_deduplication_space_saved_percent": 27,
      "volume_footprints_percent": 14
    },
    "cloud_storage": {
      "used": 402743264
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0,
      "wise_tsse_min_used_capacity_pct": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots_flexclones": {
      "logical_used": 0,
      "ratio": 0,
```

```
      "savings": 0
    },
    "footprint": 608896,
    "snapshot": {
      "available": 2000,
      "reserve_percent": 20,
      "total": 5000,
      "used": 3000,
      "used_percent": 45
    },
    "total_provisioned_space": 20971520
  },
  "state": "string",
  "uuid": "string",
  "volume-count": 0
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#aggregate[aggregate]]
a|

|warnings
|array[link:#aggregate_warning[aggregate_warning]]
a|List of validation warnings and remediation advice for the aggregate
simulate behavior.
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  },
  "num_records": 1,
  "records": [
    {
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "block_storage": {
        "hybrid_cache": {
          "disk_count": 6,
          "disk_type": "string",
          "raid_size": 24,
          "raid_type": "string",
          "simulated_raid_groups": [
            {
              "name": "string"
            }
          ],
          "size": 1612709888,
          "storage_pools": [
            {
              "storage_pool": {
                "name": "storage_pool_1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            }
          ],
          "used": 26501122
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
        "plexes": [
```

```
        {
          "name": "plex0"
        }
      ],
      "primary": {
        "checksum_style": "string",
        "disk_class": "performance",
        "disk_count": 8,
        "disk_type": "string",
        "raid_size": 16,
        "raid_type": "string",
        "simulated_raid_groups": [
          {
            "name": "string",
            "raid_type": "string"
          }
        ]
      },
      "storage_type": "string"
    },
    "create_time": "2018-01-01 12:00:00 -0400",
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "inactive_data_reporting": {
      "enabled": true,
      "start_time": "2019-12-12 12:00:00 -0400"
    },
    "inode_attributes": {
      "file_private_capacity": 31136,
      "file_public_capacity": 31136,
      "files_private_used": 502,
      "files_total": 31136,
      "files_used": 97,
      "max_files_available": 31136,
      "max_files_possible": 2844525,
      "max_files_used": 97,
      "used_percent": 5,
      "version": 4
    },
    "is_spare_low": "",
    "name": "node1_aggr_1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "snaplock_type": "string",
      "snapshot": {
        "files_total": 10,
        "files_used": 3,
        "max_files_available": 5,
        "max_files_used": 50
      },
      "space": {
        "block_storage": {
          "aggregate_metadata": 2655,
          "aggregate_metadata_percent": 8,
          "available": 10156560384,
          "data_compacted_count": 1990000,
          "data_compaction_space_saved": 1996000,
          "data_compaction_space_saved_percent": 27,
          "full_threshold_percent": 0,
          "inactive_user_data": 304448,
          "inactive_user_data_percent": 0,
          "physical_used": 2461696,
          "physical_used_percent": 50,
          "size": 10156769280,
          "used": 2088960,
          "used_including_snapshot_reserve": 674685,
          "used_including_snapshot_reserve_percent": 35,
          "used_percent": 50,
          "volume_deduplication_shared_count": 1990000,
          "volume_deduplication_space_saved": 1996000,
          "volume_deduplication_space_saved_percent": 27,
          "volume_footprints_percent": 14
        },
        "cloud_storage": {
          "used": 402743264
        },
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0,
          "wise_tsse_min_used_capacity_pct": 0
        },
        "efficiency_without_snapshots": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots_flexclones": {
```

```
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "footprint": 608896,
        "snapshot": {
          "available": 2000,
          "reserve_percent": 20,
          "total": 5000,
          "used": 3000,
          "used_percent": 45
        },
        "total_provisioned_space": 20971520
      },
      "state": "string",
      "uuid": "string",
      "volume-count": 0
    }
  ],
  "warnings": [
    {
      "action": {
        "arguments": [
          "string"
        ],
        "message": "string"
      },
      "name": "string",
      "warning": {
        "arguments": [
          "string"
        ],
        "message": "string"
      }
    }
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
```

```
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 460770
| The aggregate create job failed to create the aggregate.

| 786438
| Failed to create an aggregate on the node.

| 786439
| An aggregate already uses the specified name.

| 786446
| The node is not in cluster.

| 786468
| VLDB is offline.

| 786819
| The value is invalid for the specified option at the current privilege
level.

| 786902
| RAID-TEC aggregate is not fully supported.
```

```
| 786911
| Not every node in the cluster has the Data ONTAP version required for
the feature.

| 786973
| The number of disks specified does not meet the minimum number of disks
required for this operation.

| 787069
| Node is setup for MetroCluster over IP configuration; creating an
unmirrored aggregate is not supported in this configuration.

| 787092
| The target field cannot be specified for this operation.

| 787267
| The number of disks specified does not meet the minimum number of disks
required for the creation of a new RAID group.

| 918138
| Internal error. Failed to get encryption operation status.

| 1114292
| The required SnapLock license is not installed.

| 2425736
| No matching node found for the target UUID.

| 19726341
| Not enough eligible spare disks are available on the node.

| 19726344
| No recommendation can be made for this cluster.

| 19726357
| Automatic aggregate creation is not supported on MetroCluster with Fibre
Channel (FC).

| 19726358
| Automatic aggregate creation is not supported on ONTAP Cloud.

| 19726373
| Recommendation specified for creating aggregates is not current.

| 19726378
| Failed to create recommended aggregates on one or more nodes.
```

```
| 19726382
| Another provisioning operation is in progress on this cluster. Wait a
few minutes, and try the operation again.

| 19726386
| Encountered an error when retrieving licensing information on this
cluster.

| 19726387
| No recommendation can be provided for this cluster within the license
capacity.

| 19726401
| Aggregate recommendations are not supported when the DR group is not in
the "normal" state.

| 19726402
| Internal error. Unable to determine the MetroCluster configuration
state.

| 19726403
| Aggregate recommendation is not supported when there are no healthy
target connections to remote storage.

| 19726544
| POST on the aggregate endpoint is not supported on this version of
ONTAP.

| 65537119
| Key management is not configured for the admin SVM. Configure a key
manager for the admin SVM.

| 196608055
| Aggregate recommendation is not supported on this node because it does
not support NetApp Aggregate Encryption (NAE).

| 196608206
| Internal error. Failed to get encryption operation status.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|existing_data_disk_count
```

```
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#storage_pool_reference]
[.api-collapsible-fifth-title]
storage_pool_reference

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#storage_pools]
[.api-collapsible-fifth-title]
storage_pools

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allocation_units_count
|integer
a|Allocation count of storage pool.


|storage_pool
|link:#storage_pool_reference[storage_pool_reference]
a|Shared Storage Pool


|===


[#hybrid_cache]
[.api-collapsible-fifth-title]
hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made
up of either whole SSDs or storage pool SSDs.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|Number of disks used in the cache tier of the aggregate. Only provided
when hybrid_cache.enabled is 'true'.
```

```
|disk_type
|string
a|Type of disk being used by the aggregate's cache tier.


|enabled
|boolean
a|Specifies whether the aggregate uses HDDs with SSDs as a cache.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type for SSD cache of the aggregate. Only provided when
hybrid_cache.enabled is 'true'.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|size
|integer
a|Total usable space in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|storage_pools
|array[link:#storage_pools[storage_pools]]
a|List of storage pool properties and allocation_units_count for
aggregate.


|used
|integer
a|Space used in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|===
```

```
[#mirror]
[.api-collapsible-fifth-title]
mirror

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Aggregate is SyncMirror protected


|state
|string
a|

|===


[#plex_reference]
[.api-collapsible-fifth-title]
plex_reference

Plex


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|data_disk_count
|integer
a|Number of data disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|parity_disk_count
|integer
a|Number of parity disks in RAID group.


|raid_type
|string
```

```
a|RAID type of the aggregate.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#primary]
[.api-collapsible-fifth-title]
primary

Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|checksum_style
|string
a|The checksum style used by the aggregate.


|disk_class
|string
a|The class of disks being used by the aggregate.


|disk_count
|integer
a|Number of disks used in the aggregate. This includes parity disks, but
excludes disks in the hybrid cache.


|disk_type
|string
a|The type of disk being used by the aggregate.


|raid_size
```

```
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|hybrid_cache
|link:#hybrid_cache[hybrid_cache]
a|Contains the configuration for the hybrid cache. The hybrid cache is
made up of either whole SSDs or storage pool SSDs.


|mirror
|link:#mirror[mirror]
a|

|plexes
|array[link:#plex_reference[plex_reference]]
a|Plex reference for each plex in the aggregate.
```

```
|primary
|link:#primary[primary]
a|Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


|storage_type
|string
a|Type of aggregate.


|uses_partitions
|boolean
a|If true, aggregate is using shared disks.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

Cloud store


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#cloud_storage_tier]
[.api-collapsible-fifth-title]
cloud_storage_tier
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|cloud_store
|link:#cloud_store[cloud_store]
a|Cloud store


|used
|integer
a|Capacity used in bytes in the cloud store by this aggregate. This is a
cached value calculated every 5 minutes.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

Configuration information for the cloud storage portion of the aggregate.
This is referred to as the capacity tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|attach_eligible
|boolean
a|Specifies whether the aggregate is eligible for a cloud store to be
attached.


|migrate_threshold
|integer
a|Specifies the minimum percentage of performance tier free space that
must exist in order for migration of data from the capacity tier to
performance tier to be allowed. Only valid for PATCH operations.
```

```
|stores
|array[link:#cloud_storage_tier[cloud_storage_tier]]
a|Configuration information for each cloud storage portion of the
aggregate.


|tiering_fullness_threshold
|integer
a|The percentage of space in the performance tier that must be used before
data is tiered out to the cloud store. Only valid for PATCH operations.


|===


[#data_encryption]
[.api-collapsible-fifth-title]
data_encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drive_protection_enabled
|boolean
a|Specifies whether the aggregate uses self-encrypting drives with data
protection enabled.


|software_encryption_enabled
|boolean
a|Specifies whether NetApp aggregate encryption is enabled. All data in
the aggregate is encrypted.


|===


[#dr_home_node]
[.api-collapsible-fifth-title]
dr_home_node

Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

Node where the aggregate resides after giveback. The value for this field
might differ from the value of the 'node' field during takeover.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#inactive_data_reporting]
[.api-collapsible-fifth-title]
inactive_data_reporting
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|start_time
|string
a|Timestamp at which inactive data reporting was enabled on the aggregate.


|===


[#inode_attributes]
[.api-collapsible-fifth-title]
inode_attributes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|file_private_capacity
|integer
a|Number of files that can currently be stored on disk for system metadata
files. This number will dynamically increase as more system files are
created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|file_public_capacity
|integer
a|Number of files that can currently be stored on disk for user-visible
files.  This number will dynamically increase as more user-visible files
are created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_private_used
```

```
|integer
a|Number of system metadata files used. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_total
|integer
a|Maximum number of user-visible files that this referenced file system
can currently hold. If the referenced file system is restricted or
offline, a value of 0 is returned.


|files_used
|integer
a|Number of user-visible files used in the referenced file system. If the
referenced file system is restricted or offline, a value of 0 is returned.


|max_files_available
|integer
a|The count of the maximum number of user-visible files currently
allowable on the referenced file system.


|max_files_possible
|integer
a|The largest value to which the maxfiles-available parameter can be
increased by reconfiguration, on the referenced file system.


|max_files_used
|integer
a|The number of user-visible files currently in use on the referenced file
system.


|used_percent
|integer
a|The percentage of disk space currently in use based on user-visible file
count on the referenced file system.


|version
```

```
|integer
a|The inofile-version of the aggregate. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

The most recent sample of I/O metrics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node where the aggregate currently resides.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|files_total
|integer
a|Total files allowed in snapshots


|files_used
|integer
a|Total files created in snapshots


|max_files_available
|integer
a|Maximum files available for snapshots


|max_files_used
|integer
a|Files in use by snapshots
```

```
|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate_metadata
|integer
a|Space used by different metafiles and internal operations inside the
aggregate, in bytes.


|aggregate_metadata_percent
|integer
a|Aggregate metadata as a percentage.


|available
|integer
a|Space available in bytes.


|data_compacted_count
|integer
a|Amount of compacted data in bytes.


|data_compaction_space_saved
|integer
a|Space saved in bytes by compacting the data.


|data_compaction_space_saved_percent
|integer
a|Percentage saved by compacting the data.


|full_threshold_percent
|integer
a|The aggregate used percentage at which 'monitor.volume.full' EMS is
```

```
generated.


|inactive_user_data
|integer
a|The size that is physically used in the block storage and has a cold
temperature, in bytes. This property is only supported if the aggregate is
either attached to a cloud store or can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data or **.


|inactive_user_data_percent
|integer
a|The percentage of inactive user data in the block storage. This property
is only supported if the aggregate is either attached to a cloud store or
can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data_percent or **.


|physical_used
|integer
a|Total physical used size of an aggregate in bytes.


|physical_used_percent
|integer
a|Physical used percentage.


|size
|integer
a|Total usable space in bytes, not including WAFL reserve and aggregate
snapshot reserve.


|used
|integer
a|Space used or reserved in bytes. Includes volume guarantees and
aggregate metadata.
```

```
|used_including_snapshot_reserve
|integer
a|Total used including the snapshot reserve, in bytes.


|used_including_snapshot_reserve_percent
|integer
a|Total used including the Snapshot reserve as a percentage.


|used_percent
|integer
a|Aggregate used percentage.


|volume_deduplication_shared_count
|integer
a|Amount of shared bytes counted by storage efficiency.


|volume_deduplication_space_saved
|integer
a|Amount of space saved in bytes by storage efficiency.


|volume_deduplication_space_saved_percent
|integer
a|Percentage of space saved by storage efficiency.


|volume_footprints_percent
|integer
a|A summation of volume footprints inside the aggregate, as a percentage.
A volume's footprint is the amount of space being used for the volume in
the aggregate.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|used
|integer
a|Used space in bytes in the cloud store. Only applicable for aggregates
with a cloud store tier.


|===


[#efficiency]
[.api-collapsible-fifth-title]
efficiency

Storage efficiency.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_adaptive_compression_savings
|boolean
a|Indicates whether or not aggregate has auto adaptive compression
savings.


|cross_volume_background_dedupe
|boolean
a|Indicates whether or not cross volume background deduplication is
enabled.


|cross_volume_dedupe_savings
|boolean
a|Indicates whether or not aggregate has cross volume deduplication
savings.


|cross_volume_inline_dedupe
|boolean
```

```
a|Indicates whether or not cross volume inline deduplication is enabled.


|enable_workload_informed_tsse
|boolean
a|Indicates whether Workload Informed TSSE is enabled on the system.


|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|wise_tsse_min_used_capacity_pct
|integer
a|Minimum amount of used data in aggregate required to trigger cold
compression on TSSE volumes.


|===


[#efficiency_without_snapshots]
[.api-collapsible-fifth-title]
efficiency_without_snapshots

Storage efficiency that does not include the savings provided by
snapshots.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
```

```
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots_flexclones]
[.api-collapsible-fifth-title]
efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by snapshots
and Flexclone volumes.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===
```

```
[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|Available space for snapshots in bytes


|reserve_percent
|integer
a|Percentage of space reserved for snapshots


|total
|integer
a|Total space for snapshots in bytes


|used
|integer
a|Space used by snapshots in bytes


|used_percent
|integer
a|Percentage of disk space used by snapshots


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|block_storage
|link:#block_storage[block_storage]
a|

|cloud_storage
|link:#cloud_storage[cloud_storage]
a|

|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and Flexclone volumes.


|footprint
|integer
a|A summation of volume footprints (including volume guarantees), in
bytes. This includes all of the volume footprints in the block_storage
tier and the cloud_storage tier.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|snapshot
|link:#snapshot[snapshot]
a|

|total_provisioned_space
|integer
a|A summation of volume size, in bytes. This includes all of the volumes
```

```
inside the aggregate.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either total_provisioned_space or **.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
```

```
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

The real time I/O statistics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
```

```
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|

|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|

|inode_attributes
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
```

```
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snaplock_type
|string
a|SnapLock type.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
|string
a|Operational state of the aggregate.


|uuid
|string
a|Aggregate UUID.


|volume-count
|integer
a|Number of volumes in the aggregate.
```

```
|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#action]
[.api-collapsible-fifth-title]
action

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the specified action message.


|code
|integer
a|Corrective action code of the specified action.


|message
|string
a|Specifies the corrective action to be taken to resolve the issue.
```

```
|===


[#warning]
[.api-collapsible-fifth-title]
warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the warning message encountered.


|code
|integer
a|Warning code of the warning encountered.


|message
|string
a|Details of the warning encountered by the aggregate simulate query.


|===


[#aggregate_warning]
[.api-collapsible-fifth-title]
aggregate_warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
|link:#action[action]
a|

|name
```

```
|string
a|Name of the entity that returns the warning.


|warning
|link:#warning[warning]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


[[ID602d816babea8b65ad97de0e3031aa15]]
= Retrieve a collection of cloud stores used by an aggregate

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/cloud-stores`#

*Introduced In:* 9.6

Retrieves the collection of cloud stores used by an aggregate.

== Related ONTAP commands

* `storage aggregate object-store show`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|used
|integer
|query
|False
a|Filter by used


|availability_at_partner
|string
|query
|False
a|Filter by availability_at_partner

* Introduced in: 9.15


|unavailable_reason.message
|string
|query
|False
a|Filter by unavailable_reason.message

* Introduced in: 9.7


|availability
|string
|query
|False
a|Filter by availability


|target.uuid
```

```
|string
|query
|False
a|Filter by target.uuid


|target.name
|string
|query
|False
a|Filter by target.name


|primary
|boolean
|query
|False
a|Filter by primary


|resync-progress
|integer
|query
|False
a|Filter by resync-progress

* Introduced in: 9.14


|mirror_degraded
|boolean
|query
|False
a|Filter by mirror_degraded


|aggregate.name
|string
|query
|False
a|Filter by aggregate.name

* Introduced in: 9.9


|unreclaimed_space_threshold
|integer
```

```
|query
|False
a|Filter by unreclaimed_space_threshold


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
```

```
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#cloud_store[cloud_store]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
```

```
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregate": {
        "name": "aggr1"
      },
      "availability": "string",
      "availability_at_partner": "string",
      "resync-progress": 0,
      "target": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "target1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "unavailable_reason": {
        "message": "string"
      },
      "unreclaimed_space_threshold": 20,
      "used": 0
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|

|===


[#target]
[.api-collapsible-fifth-title]
target

Cloud target


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#unavailable_reason]
[.api-collapsible-fifth-title]
unavailable_reason

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|message
|string
a|Indicates why the object store is unavailable.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|aggregate
|link:#aggregate[aggregate]
a|Aggregate


|availability
|string
a|Availability of the object store.


|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
primary cloud store of a FabricPool.


|primary
|boolean
```

```
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.


|target
|link:#target[target]
a|Cloud target


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|

|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID3b052066e077aa73178b6902c14ed1dd]]
= Attach an object store to an aggregate or add a second object store as a
mirror

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/cloud-stores`#

*Introduced In:* 9.6

Attaches an object store to an aggregate, or adds a second object store as
a mirror.

== Required properties

* `target.uuid` or `target.name` - UUID or name of the cloud target.

== Recommended optional properties

* `primary` - _true_ if the object store is primary or _false_ if it is a
mirror.
* `allow_flexgroups` - Allow attaching object store to an aggregate
containing FlexGroup constituents.
* `check_only` - Validate only and do not add the cloud store.

== Default property values

* `primary` - _true_
* `allow_flexgroups` - _false_
* `check_only` - _false_

== Related ONTAP commands

* `storage aggregate object-store attach`
* `storage aggregate object-store mirror`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|allow_flexgroups
|boolean
|query
|False
a|This optional parameter allows attaching object store to an aggregate
containing FlexGroup constituents. The default value is false. Mixing
FabricPools and non-FabricPools within a FlexGroup is not recommended. All
aggregates hosting constituents of a FlexGroup should be attached to the
object store.


|check_only
|boolean
|query
|False
a|Validate only and do not add the cloud store.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
```

```
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|Aggregate


|availability
|string
a|Availability of the object store.


|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
primary cloud store of a FabricPool.


|primary
|boolean
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.
```

```
|target
|link:#target[target]
a|Cloud target


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|

|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "aggregate": {
    "name": "aggr1"
  },
  "availability": "string",
  "availability_at_partner": "string",
  "resync-progress": 0,
  "target": {
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

```
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 787143
| Object store is already attached to FabricPool.

| 787144
| Aggregate is not a FabricPool.

| 787205
| The operation failed because the previous unmirror operation on the
aggregate is still in progress. Check that the previous unmirror operation
is complete and try the command again.

| 787218
| The specified object store does not exist.

| 7209876
| Aggregate is already a FabricPool Aggregate.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#target]
[.api-collapsible-fifth-title]
target

Cloud target


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|

|uuid
|string
a|

|===


[#unavailable_reason]
[.api-collapsible-fifth-title]
unavailable_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|message
|string
a|Indicates why the object store is unavailable.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|Aggregate


|availability
|string
a|Availability of the object store.
```

```
|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
primary cloud store of a FabricPool.


|primary
|boolean
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.


|target
|link:#target[target]
a|Cloud target


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|

|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDa1eca0bcbec6e0445852ec196273a574]]
= Remove a cloud target from an aggregate

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}`#

*Introduced In:* 9.6

Removes the specified cloud target from the aggregate. Only removal of a
mirror is allowed. The primary cannot be removed. This request starts a
job and returns a link to that job.

== Related ONTAP commands
```

```
* `storage aggregate object-store unmirror`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|target.uuid
|string
|path
|True
a|Cloud target UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 3
| Invalid operation
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID96e8c905c4c105bf013e2b5126b03ec4]]
= Retrieve the cloud store for an aggregate

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}`#

*Introduced In:* 9.6

Retrieves the cloud store for the aggregate using the specified cloud
target UUID.

== Related ONTAP commands

* `storage aggregate object-store show`


== Parameters

[cols=5*,options=header]
```

```
|===

|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|target.uuid
|string
|path
|True
a|Cloud target UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|aggregate
|link:#aggregate[aggregate]
```

```
a|Aggregate


|availability
|string
a|Availability of the object store.


|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
primary cloud store of a FabricPool.


|primary
|boolean
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.


|target
|link:#target[target]
a|Cloud target


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|

|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
```

```
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregate": {
    "name": "aggr1"
  },
  "availability": "string",
  "availability_at_partner": "string",
  "resync-progress": 0,
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#target]
```

```
[.api-collapsible-fifth-title]
target

Cloud target


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#unavailable_reason]
[.api-collapsible-fifth-title]
unavailable_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|message
|string
a|Indicates why the object store is unavailable.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID215bfb56a80891ade12423a988f5ae20]]
= Update a cloud store

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}`#

*Introduced In:* 9.6

Updates the cloud store specified by the UUID with the fields in the body.
This request starts a job and returns a link to that job.

== Related ONTAP commands

* `storage aggregate object-store modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|target.uuid
```

```
|string
|path
|True
a|Cloud target UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|availability
|string
a|Availability of the object store.


|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
```

```
primary cloud store of a FabricPool.


|primary
|boolean
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|

|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "availability": "string",
  "availability_at_partner": "string",
  "resync-progress": 0,
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

```
====
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 787154
```

```
| The Object stores are currently not synchronized. Check the mirror
resync status.

| 787156
| Modifying the attributes of a mirror object store is not allowed.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#target]
[.api-collapsible-fifth-title]
target

Cloud target


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#unavailable_reason]
[.api-collapsible-fifth-title]
unavailable_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|message
```

```
|string
a|Indicates why the object store is unavailable.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

[cols=3*,options=header]
|===
|Name
|Type
|Description

|availability
|string
a|Availability of the object store.


|availability_at_partner
|string
a|Availability of the object store at the HA partner.


|mirror_degraded
|boolean
a|This field identifies if the mirror cloud store is in sync with the
primary cloud store of a FabricPool.


|primary
|boolean
a|This field indicates whether the cloud store is the primary cloud store
of a mirrored FabricPool.


|resync-progress
|integer
a|Resync progress of the mirror object store in percentage.


|unavailable_reason
|link:#unavailable_reason[unavailable_reason]
a|
```

```
|unreclaimed_space_threshold
|integer
a|Usage threshold for reclaiming unused space in the cloud store. Valid
values are 0 to 99. The default value depends on the provider type. This
can be specified in PATCH but not POST.


|used
|integer
a|The amount of object space used. Calculated every 5 minutes and cached.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
```

```
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
= Manage storage aggregate plexes

:leveloffset: +1


[[ID2ae2aaedd5c74714f45fd9480e851f1c]]
= Manage storage aggregate plexes



== Overview

The Storage Aggregate Plex API provides relevant state information for
each plex in the aggregate.
For each plex, details are provided for the RAID groups in the plex and
the disks that make up each RAID group.

== Examples

=== Retrieving all aggregates and plexes

The following example shows the response with a list of aggregates and
plexes:

----

# the API:
/storage/aggregates/{aggregate.uuid}/plexes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/*/plexes" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "aggregate": {
      "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29",
      "name": "test2"
    },
    "name": "plex0"
  },
  {
    "aggregate": {
```

```
      "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29",
      "name": "test2"
    },
    "name": "plex1",
  },
  {
    "aggregate": {
      "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af",
      "name": "test"
    },
    "name": "plex0"
  },
  {
    "aggregate": {
      "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af",
      "name": "test"
    },
    "name": "plex1"
  },
  {
    "aggregate": {
      "uuid": "7ee89e48-5d81-4609-9e1b-5d8d0995a886",
      "name": "aggr1",
    },
    "name": "plex0",
  },
  {
    "aggregate": {
      "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40",
      "name": "test3",
    },
    "name": "plex0",
  },
  {
    "aggregate": {
      "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40",
      "name": "test3",
    },
    "name": "plex1"
  },
  {
    "aggregate": {
      "uuid": "8f13de5c-99cf-4ada-884c-3cc32deb304a",
      "name": "aggr2",
    },
    "name": "plex0"
```

```
  }
],
"num_records": 8
}
----

=== Retrieving a specific plex in all aggregates

The following example shows the response with a list of specific plexes in
all aggregates:

----

# the API:
/storage/aggregates/{aggregate.uuid}/plexes/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/*/plexes/plex0" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "aggregate": {
      "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29",
      "name": "test2"
    },
    "name": "plex0"
  },
  {
    "aggregate": {
      "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af",
      "name": "test"
    },
    "name": "plex0"
  },
  {
    "aggregate": {
      "uuid": "7ee89e48-5d81-4609-9e1b-5d8d0995a886",
      "name": "aggr1",
    },
    "name": "plex0",
  },
  {
    "aggregate": {
```

```
      "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40",
      "name": "test3",
    },
    "name": "plex0",
  },
  {
    "aggregate": {
      "uuid": "8f13de5c-99cf-4ada-884c-3cc32deb304a",
      "name": "aggr2",
    },
    "name": "plex0"
  }
],
"num_records": 8
}
----

=== Retrieving the list of plexes in an aggregate

The following example shows the response with the list of plexes in an
aggregate:

----

# The API:
/api/storage/aggregates/{uuid}/plexes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-
8f01-712f626c983c/plexes" -H "accept: application/json"

# The response:
{
"records": [
  {
    "name": "plex0",
  },
  {
    "name": "plex4",
  }
],
"num_records": 2
}
----

=== Retrieving a specific plex in an aggregate
```

```
The following example shows the response when requesting a specific plex
of an aggregate:

----

# The API:
/api/storage/aggregates/{uuid}/plexes/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-
8f01-712f626c983c/plexes/plex0" -H "accept: application/json"

# The response:
{
"aggregate": {
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
},
"name": "plex0",
"online": true,
"state": "normal",
"pool": "pool0",
"resync": {
  "active": false
},
"raid_groups": [
  {
    "name": "rg0",
    "raid_type": "raid_dp",
    "cache_tier": false,
    "degraded": false,
    "recomputing_parity": {
      "active": false
    },
    "reconstruct": {
      "active": false
    },
    "disks": [
      {
        "position": "dparity",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.29",
        }
```

```
      },
      {
        "position": "parity",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.4",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.30",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.5",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.31",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.6",
        }
      }
```

```
    ]
  }
]
}
----
```

```
[[ID2efbe21cf07bd4691d249a808371dd04]]
= Retrieve a collection of plexes for an aggregate

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/plexes`#

*Introduced In:* 9.6

Retrieves the collection of plexes for the specified aggregate.

== Related ONTAP commands

* `storage aggregate plex show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
a|Aggregate UUID


|online
|boolean
|query
|False
```

```
a|Filter by online


|name
|string
|query
|False
a|Filter by name


|state
|string
|query
|False
a|Filter by state


|resync.percent
|integer
|query
|False
a|Filter by resync.percent


|resync.active
|boolean
|query
|False
a|Filter by resync.active


|resync.level
|string
|query
|False
a|Filter by resync.level


|pool
|string
|query
|False
a|Filter by pool


|raid_groups.raid_type
|string
```

```
|query
|False
a|Filter by raid_groups.raid_type

* Introduced in: 9.9


|raid_groups.reconstruct.active
|boolean
|query
|False
a|Filter by raid_groups.reconstruct.active

* Introduced in: 9.7


|raid_groups.reconstruct.percent
|integer
|query
|False
a|Filter by raid_groups.reconstruct.percent

* Introduced in: 9.7


|raid_groups.disks.disk.name
|string
|query
|False
a|Filter by raid_groups.disks.disk.name

* Introduced in: 9.7


|raid_groups.disks.usable_size
|integer
|query
|False
a|Filter by raid_groups.disks.usable_size

* Introduced in: 9.7


|raid_groups.disks.type
|string
|query
|False
```

```
a|Filter by raid_groups.disks.type

* Introduced in: 9.7


|raid_groups.disks.state
|string
|query
|False
a|Filter by raid_groups.disks.state

* Introduced in: 9.7


|raid_groups.disks.position
|string
|query
|False
a|Filter by raid_groups.disks.position

* Introduced in: 9.7


|raid_groups.cache_tier
|boolean
|query
|False
a|Filter by raid_groups.cache_tier

* Introduced in: 9.7


|raid_groups.recomputing_parity.percent
|integer
|query
|False
a|Filter by raid_groups.recomputing_parity.percent

* Introduced in: 9.7


|raid_groups.recomputing_parity.active
|boolean
|query
|False
a|Filter by raid_groups.recomputing_parity.active
```

```
* Introduced in: 9.7


|raid_groups.name
|string
|query
|False
a|Filter by raid_groups.name

* Introduced in: 9.7


|raid_groups.degraded
|boolean
|query
|False
a|Filter by raid_groups.degraded

* Introduced in: 9.7


|aggregate.name
|string
|query
|False
a|Filter by aggregate.name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
```

```
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#plex[plex]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
      },
      "name": "plex0",
      "pool": "string",
      "raid_groups": [
        {
          "disks": [
            {
              "disk": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "1.0.1"
              },
              "position": "string",
              "state": "string",
              "type": "ssd",
              "usable_size": 947912704
            }
          ],
          "name": "rg0",
          "raid_type": "raid_dp",
          "recomputing_parity": {
            "percent": 10
          },
          "reconstruct": {
            "percent": 10
          }
        }
      ],
      "resync": {
        "level": "string",
        "percent": 10
      },
      "state": "string"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
```

```
a|

|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#disk]
[.api-collapsible-fifth-title]
disk

Disk


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|name
|string
a|

|===


[#raid_group_disk]
[.api-collapsible-fifth-title]
raid_group_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Disk


|position
|string
a|The position of the disk within the RAID group.


|state
|string
a|The state of the disk within the RAID group.


|type
|string
a|Disk interface type


|usable_size
|integer
a|Size in bytes that is usable by the aggregate.


|===
```

```
[#recomputing_parity]
[.api-collapsible-fifth-title]
recomputing_parity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|RAID group is recomputing parity


|percent
|integer
a|Recomputing parity percentage


|===


[#reconstruct]
[.api-collapsible-fifth-title]
reconstruct

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|One or more disks in this RAID group are being reconstructed.


|percent
|integer
a|Reconstruct percentage


|===
```

```
[#raid_group]
[.api-collapsible-fifth-title]
raid_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cache_tier
|boolean
a|RAID group is a cache tier


|degraded
|boolean
a|RAID group is degraded. A RAID group is degraded when at least one disk
from that group has failed or is offline.


|disks
|array[link:#raid_group_disk[raid_group_disk]]
a|

|name
|string
a|RAID group name


|raid_type
|string
a|RAID type of the raid group.


|recomputing_parity
|link:#recomputing_parity[recomputing_parity]
a|

|reconstruct
|link:#reconstruct[reconstruct]
a|

|===


[#resync]
```

```
[.api-collapsible-fifth-title]
resync

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|Plex is being resynchronized to its mirrored plex


|level
|string
a|Plex resyncing level


|percent
|integer
a|Plex resyncing percentage


|===


[#plex]
[.api-collapsible-fifth-title]
plex

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|Aggregate


|name
|string
a|Plex name
```

```
|online
|boolean
a|Plex is online


|pool
|string
a|SyncMirror pool assignment


|raid_groups
|array[link:#raid_group[raid_group]]
a|

|resync
|link:#resync[resync]
a|

|state
|string
a|Plex state


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDfff11cc75ff8f89b6048c9956b9edd48]]
= Retrieve a plex specified by the aggregate UUID and plex name

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{aggregate.uuid}/plexes/{name}`#

*Introduced In:* 9.6

Retrieves the plex specified by the aggregate UUID and plex name.

== Related ONTAP commands

* `storage aggregate plex show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|aggregate.uuid
|string
|path
|True
```

```
a|Aggregate UUID


|name
|string
|path
|True
a|Plex name


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|Aggregate


|name
|string
a|Plex name


|online
|boolean
a|Plex is online


|pool
|string
a|SyncMirror pool assignment
```

```
|raid_groups
|array[link:#raid_group[raid_group]]
a|

|resync
|link:#resync[resync]
a|

|state
|string
a|Plex state


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "plex0",
  "pool": "string",
  "raid_groups": [
    {
      "disks": [
        {
          "disk": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "1.0.1"
          },
```

```
          "position": "string",
          "state": "string",
          "type": "ssd",
          "usable_size": 947912704
        }
      ],
      "name": "rg0",
      "raid_type": "raid_dp",
      "recomputing_parity": {
        "percent": 10
      },
      "reconstruct": {
        "percent": 10
      }
    }
  ],
  "resync": {
    "level": "string",
    "percent": 10
  },
  "state": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#disk]
[.api-collapsible-fifth-title]
disk

Disk


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===


[#raid_group_disk]
[.api-collapsible-fifth-title]
raid_group_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Disk


|position
|string
a|The position of the disk within the RAID group.


|state
|string
a|The state of the disk within the RAID group.


|type
|string
a|Disk interface type


|usable_size
|integer
```

```
a|Size in bytes that is usable by the aggregate.


|===


[#recomputing_parity]
[.api-collapsible-fifth-title]
recomputing_parity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|RAID group is recomputing parity


|percent
|integer
a|Recomputing parity percentage


|===


[#reconstruct]
[.api-collapsible-fifth-title]
reconstruct

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|One or more disks in this RAID group are being reconstructed.


|percent
|integer
a|Reconstruct percentage
```

```
|===


[#raid_group]
[.api-collapsible-fifth-title]
raid_group

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cache_tier
|boolean
a|RAID group is a cache tier


|degraded
|boolean
a|RAID group is degraded. A RAID group is degraded when at least one disk
from that group has failed or is offline.


|disks
|array[link:#raid_group_disk[raid_group_disk]]
a|

|name
|string
a|RAID group name


|raid_type
|string
a|RAID type of the raid group.


|recomputing_parity
|link:#recomputing_parity[recomputing_parity]
a|

|reconstruct
|link:#reconstruct[reconstruct]
a|
```

```
|===


[#resync]
[.api-collapsible-fifth-title]
resync

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active
|boolean
a|Plex is being resynchronized to its mirrored plex


|level
|string
a|Plex resyncing level


|percent
|integer
a|Plex resyncing percentage


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1
```

```
= Manage storage aggregates

:leveloffset: +1


[[ID908f3ab78d5982a119fa56a43ad0a391]]
= Manage storage aggregates



== Updating storage aggregates

The PATCH operation is used to modify properties of the aggregate. There
are several properties that can be modified on an aggregate. Only one
property can be modified for each PATCH request.
PATCH operations on the aggregate's disk count will be blocked while one
or more nodes in the cluster are simulating or implementing automatic
aggregate creation.

Aggregates are automatically managed in the ASA r2 system. Therefore,
PATCH and DELETE operations are not supported in ASA r2.

The following is a list of properties that can be modified using the PATCH
operation including a brief description for each:

* name - This property can be changed to rename the aggregate.
* node.name and node.uuid - Either property can be updated in order to
relocate the aggregate to a different node in the cluster.
* state - This property can be changed to 'online' or 'offline'. Setting
an aggregate 'offline' would automatically offline all the volumes
currently hosted on the aggregate.
* block_storage.mirror.enabled - This property can be changed from 'false'
to 'true' in order to mirror the aggregate, if the system is capable of
doing so.
* block_storage.primary.disk_count - This property can be updated to
increase the number of disks in an aggregate.
* block_storage.primary.raid_size - This property can be updated to set
the desired RAID size.
* block_storage.primary.raid_type - This property can be updated to set
the desired RAID type.
* cloud_storage.tiering_fullness_threshold - This property can be updated
to set the desired tiering fullness threshold if using FabricPool.
* cloud_storage.migrate_threshold - This property can be updated to set
the desired migrate threshold if using FabricPool.
* data_encryption.software_encryption_enabled - This property enables or
```

```
disables NAE on the aggregate.
* block_storage.hybrid_cache.storage_pools.allocation_units_count - This
property can be updated to add a storage pool to the aggregate specifying
the number of allocation units.
* block_storage.hybrid_cache.storage_pools.name - This property can be
updated to add a storage pool to the aggregate specifying the storage pool
name. block_storage.hybrid_cache.storage_pools.uuid or this field must be
specified with
block_storage.hybrid_cache.storage_pools.allocation_units_count.
* block_storage.hybrid_cache.storage_pools.uuid - This property can be
updated to add a storage pool to the aggregate specifying the storage pool
uuid. block_storage.hybrid_cache.storage_pools.name or this field must be
specified with
block_storage.hybrid_cache.storage_pools.allocation_units_count.
* block_storage.hybrid_cache.raid_size - This property can be updated to
set the desired RAID size. This property can also be specified on the
first time addition of a storage pool to the aggregate.
* block_storage.hybrid_cache.raid_type - This property can be updated to
set the desired RAID type of a physical SSD Flash Pool. This property can
also be specified on the first time addition of a storage pool to the
aggregate. When specifying a raidtype of raid4, the node is required to
have spare SSDs for the storage pool as well.
* block_storage.hybrid_cache.disk_count - This property can be specified
on the first time addition of physical SSD cache to the aggregate. It can
also be updated to increase the number of disks in the physical SSD cache
of a hybrid aggregate.

=== Aggregate expansion

The PATCH operation also supports automatically expanding an aggregate
based on the spare disks which are present within the system. Running
PATCH with the query "auto_provision_policy" set to "expand" starts the
recommended expansion job. In order to see the expected change in capacity
before starting the job, call GET on an aggregate instance with the query
"auto_provision_policy" set to "expand".

=== Manual simulated aggregate expansion

The PATCH operation also supports simulated manual expansion of an
aggregate.
Running PATCH with the query "simulate" set to "true" and
"block_storage.primary.disk_count" set to the final disk count will start
running the prechecks associated with expanding the aggregate to the
proposed size.
The response body will include information on how many disks the aggregate
can be expanded to, any associated warnings, along with the proposed final
```

```
size of the aggregate.

== Deleting storage aggregates

If volumes exist on an aggregate, they must be deleted or moved before the
aggregate can be deleted.
See the /storage/volumes API for details on moving or deleting volumes.

== Adding a storage pool to an aggregate

A storage pool can be added to an aggregate by patching the field
"block_storage.hybrid_cache.storage_pools.allocation_units_count" while
also specifying the specific storage pool using the
"block_storage.hybrid_cache.storage_pools.name" or
"block_storage.hybrid_cache.storage_pools.uuid". Subsequent patches to the
aggregate can be completed to increase allocation unit counts
or adding additional storage pools. On the first time addition of a
storage pool to the aggregate, the raidtype can be optionally specified
using the "block_storage.hybrid_cache.raid_type" field.

== Adding physical SSD cache capacity to an aggregate

The PATCH operation supports addition of a new physical SSD cache to an
aggregate. It also supports expansion of existing physical SSD cache in
the hybrid aggregate.
Running PATCH with "block_storage.hybrid_cache.disk_count" set to the
final disk count will expand the physical SSD cache of the hybrid
aggregate to the proposed size.
The RAID type can be optionally specified using the
"block_storage.hybrid_cache.raid_type" field.
The RAID size can be optionally specified using the
"block_storage.hybrid_cache.raid_size" field.
These operations can also be simulated by setting the query "simulate" to
"true".

'''

== Examples

=== Retrieving a specific aggregate from the cluster

The following example shows the response of the requested aggregate. If
there is no aggregate with the requested UUID, an error is returned.

----

# The API:
```

```
/api/storage/aggregates/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/870dd9f2-bdfa-4167-
b692-57d1cec874d4" -H "accept: application/json"

# The response:
{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "test1",
"node": {
  "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
  "name": "node-1",
},
"home_node": {
  "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
  "name": "node-1",
},
"space": {
  "block_storage": {
    "size": 235003904,
    "available": 191942656,
    "used": 43061248,
    "full_threshold_percent": 98,
    "physical_used": 5271552,
    "physical_used_percent": 1,
    "volume_footprints_percent": 14,
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "data_compacted_count": 666666,
    "data_compaction_space_saved": 654566,
    "data_compaction_space_saved_percent": 47,
    "volume_deduplication_shared_count": 567543,
    "volume_deduplication_space_saved": 23765,
    "volume_deduplication_space_saved_percent": 32,
    "used_percent": 50,
  },
  "snapshot": {
    "used_percent": 45,
    "available": 2000,
    "total": 5000,
    "used": 3000,
    "reserve_percent": 20
  },
```

```
    "cloud_storage": {
      "used": 0
    },
    "efficiency": {
      "savings": 1408029,
      "ratio": 6.908119720880661,
      "logical_used": 1646350,
      "cross_volume_background_dedupe": true,
      "cross_volume_inline_dedupe": false,
      "cross_volume_dedupe_savings": true,
      "auto_adaptive_compression_savings": false,
      "enable_workload_informed_tsse": true,
      "wise_tsse_min_used_capacity_pct": 2
    },
    "efficiency_without_snapshots": {
      "savings": 0,
      "ratio": 1,
      "logical_used": 737280
    },
    "efficiency_without_snapshots_flexclones": {
      "savings": 5000,
      "ratio": 2,
      "logical_used": 10000
    }
  },
  "snapshot": {
    "files_total": 10,
    "files_used": 3,
    "max_files_available": 5,
    "max_files_used": 50
  },
  "state": "online",
  "snaplock_type": "non_snaplock",
  "create_time": "2018-12-04T15:40:38-05:00",
  "data_encryption": {
    "software_encryption_enabled": false,
    "drive_protection_enabled": false
  },
  "block_storage": {
    "uses_partitions": false,
    "storage_type": "vmdisk",
    "primary": {
      "disk_count": 6,
      "disk_class": "solid_state",
      "raid_type": "raid_dp",
      "raid_size": 24,
```

```
    "checksum_style": "block",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false,
    "state": "unmirrored"
  },
  "plexes": [
    {
      "name": "plex0",
    }
  ]
},
"cloud_storage": {
  "attach_eligible": false
},
"inode_attributes": {
  "files_total": 31136,
  "files_used": 97,
  "max_files_available": 31136,
  "max_files_possible": 2844525,
  "max_files_used": 97,
  "used_percent": 5
},
"volume_count": 0,
}
----

=== Retrieving statistics and metric for an aggregate

In this example, the API returns the "statistics" and "metric" properties
for the aggregate requested.

----

#The API:
/api/storage/aggregates/{uuid}?fields=statistics,metric

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-
bad0-005056b48388?fields=statistics,metric" -H "accept: application/json"

#The response:
```

```
{
"uuid": "538bf337-1b2c-11e8-bad0-005056b48388",
"name": "aggr4",
"metric": {
     "timestamp": "2019-07-08T22:16:45Z",
     "duration": "PT15S",
     "status": "ok",
     "throughput": {
       "read": 7099,
       "write": 840226,
       "other": 193293789,
       "total": 194141115
       },
     "latency": {
       "read": 149,
       "write": 230,
       "other": 123,
       "total": 124
     },
     "iops": {
       "read": 1,
       "write": 17,
       "other": 11663,
       "total": 11682
     },
 },
  "statistics": {
     "timestamp": "2019-07-08T22:17:09Z",
     "status": "ok",
     "throughput_raw": {
       "read": 3106045952,
       "write": 63771742208,
       "other": 146185560064,
       "total": 213063348224
     },
     "latency_raw": {
       "read": 54072313,
       "write": 313354426,
       "other": 477201985,
       "total": 844628724
     },
     "iops_raw": {
       "read": 328267,
       "write": 1137230,
       "other": 1586535,
       "total": 3052032
```

```
      }
  },
}
----

For more information and examples on viewing historical performance
metrics for any given aggregate, see
xref:{relative_path}storage_aggregates_uuid_metrics_endpoint_overview.html
l[DOC /storage/aggregates/\{uuid}/metrics]

=== Simulating aggregate expansion

The following example shows the response for a simulated data aggregate
expansion based on the values of the 'block_storage.primary.disk_count'
attribute passed in.
The query does not modify the existing aggregate but returns how the
aggregate will look after the expansion along with any associated
warnings.
Simulated data aggregate expansion will be blocked while one or more nodes
in the cluster are simulating or implementing automatic aggregate
creation.
This will be reflected in the following attributes:

* space.block_storage.size - Total usable space in bytes, not including
WAFL reserve and aggregate snapshot reserve.
* block_storage.primary.disk_count - Number of disks that could be used to
create the aggregate.

----

# The API:
/api/storage/aggregates/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-
42bd-babb-ef437d118314?simulate=true" -H "accept: application/json" -d
"{\"block_storage\": {\"primary\": {\"disk_count\": 13}}}"

# The response:
{
"warnings": [
  {
    "name": "node_2_SSD_1",
    "warning": {
      "message": "One or more disks will not be added. 10 disks specified,
9 disks will be added.",
```

```
      "code": 787170,
      "arguments": [
        "10",
        "9"
      ]
    }
  }
],
"num_records": 1,
"records": [
  {
    "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
    "name": "node_2_SSD_1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2"
    },
    "space": {
      "block_storage": {
        "size": 1116180480
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 12,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 12,
        "simulated_raid_groups": [
           {
             "name": "test/plex0/rg0",
             "existing_parity_disk_count": 2,
             "added_parity_disk_count": 0,
             "existing_data_disk_count": 1,
             "added_data_disk_count": 9,
             "usable_size": 12309487,
             "is_partition": false
           }
         ]
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": false
```

```
      }
    },
  }
]
}
----

=== Manual aggregate expansion with disk size query

The following example shows the response for aggregate expansion based on
the values of the 'block_storage.hybrid_cache.disk_count' attribute based
on the disk size passed in.

----

# The API:
/api/storage/aggregate/{uuid}?disk_size={disk_size}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-
42bd-babb-ef437d118314?disk_size=1902379008" -H "accept: application/json"
-d "{\"block_storage\": {\"hybrid_cache\": {\"disk_count\": 4}}}"

# The response:
{
"job": {
  "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
    }
  }
}
}
----

=== Simulating a manual aggregate expansion with disk size query

The following example shows the response for a manual aggregate expansion
based on the values of the 'block_storage.hybrid_cache.disk_count'
attribute based on the disk size passed in.
The query internally maps out the appropriate expansion as well as
warnings that may be associated for the hybrid enabled aggregate.

----

# The API:
```

```
/api/storage/aggregate/{uuid}?simulate=true&disk_size=1902379008

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-
42bd-babb-ef437d118314?simulate=true&disk_size=1902379008" -H "accept:
application/json" -d "{\"block_storage\": {\"hybrid_cache\":
{\"disk_count\": 4}}}"

# The response:
{
"num_records": 1,
"records": [
  {
    "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
    "name": "ag1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-
005056bb2030"
        }
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 4,
        "disk_class": "virtual",
        "raid_type": "raid_dp",
        "disk_type": "vm_disk",
      },
      "hybrid_cache": {
        "disk_type": "ssd",
        "enabled": true,
        "disk_count": 4,
        "raid_type": "raid_dp",
        "size": 3761766400,
        "simulated_raid_groups": [
          {
            "name": "test/plex0/rg0",
            "existing_parity_disk_count": 2,
            "existing_data_disk_count": 1,
            "added_parity_disk_count": 0,
            "added_data_disk_count": 1,
            "usable_size": 1880883200,
```

```
            "is_partition": false
          },
        ]
      },
      "mirror": {
        "enabled": false
      },
      "_links": {
        "self": {
          "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-
ef437d118314"
        }
      }
    }
  }
]
}
----

=== Simulating a manual aggregate expansion with raid group query

The following example shows the response for a manual aggregate expansion
based on the values of the 'block_storage.primary.disk_count' attribute
passed in.
The query internally maps out the appropriate expansion as well as
warnings that may be associated and lays out the new raidgroups in a more
detailed view. An additional query can be passed in to specify
raidgroup addition by new raidgroup, all raidgroups or a specific
raidgroup.

----

# The API:
/api/storage/aggregate/{uuid}?simulate=true&raid_group=[new&#124;all&#124;
rgX]

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-
42bd-babb-ef437d118314?simulate=true&raid_group=new" -H "accept:
application/json" -d "{\"block_storage\": {\"primary\": {\"disk_count\":
24}}}"

# The response:
{
"warnings": [
  {
```

```
    "name": "test",
    "warning": {
      "code": 11,
      "message": "Number of unassigned disks attached to node \"node-2\":
6.",
      "arguments": [
        "6",
        "node-2"
      ]
    }
  }
],
"num_records": 1,
"records": [
  {
    "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
    "name": "test",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2"
    },
    "space": {
      "block_storage": {
        "size": 33292025856
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 24,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 24,
        "simulated_raid_groups": [
           {
             "name": "test/plex0/rg0",
             "existing_parity_disk_count": 0,
             "added_parity_disk_count": 2,
             "existing_data_disk_count": 0,
             "added_data_disk_count": 10,
             "usable_size": 12309487,
             "is_partition": false
           },
           {
             "name": "test/plex1/rg1",
             "existing_parity_disk_count": 0,
```

```
              "added_parity_disk_count": 2,
              "existing_data_disk_count": 0,
              "added_data_disk_count": 10,
              "usable_size": 12309487,
              "is_partition": false
            }
          ]
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": false
      }
    }
  }
]
}
----

=== Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare
information for the expansion of this particular aggregate. The output is
restricted to only spares that are compatible with this aggregate.

----

# The API:
/api/storage/aggregates?show_spares=true&uuid={uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates?uuid=cae60cfe-deae-
42bd-babb-ef437d118314&show_spares=true" -H "accept: application/json"

# The response:
{
"records": [],
"num_records": 0,
"spares": [
  {
    "node": {
      "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
      "name": "node-2"
    },
    "disk_class": "solid_state",
```

```
    "disk_type": "ssd",
    "size": 3720609792,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "usable": 12,
    "layout_requirements": [
      {
        "raid_type": "raid_dp",
        "default": true,
        "aggregate_min_disks": 3,
        "raid_group": {
          "min": 3,
          "max": 28,
          "default": 24
        }
      }
    ]
  }
]
}
----

=== Retrieving the SSD spare count for the cluster

The following example shows the response from retrieving SSD spare count
information for the expansion of this particular aggregate's hybrid cache
tier. The output is restricted to only spares that are compatible with
this aggregate.

----

# The API:
/api/storage/aggregates?show_spares=true&uuid={uuid}&flash_pool_eligible=t
rue

# The response:
{
"records": [],
"num_records": 0,
"spares": [
  {
    "node": {
      "uuid": "c35c5975-cbcb-11ec-a3e1-005056bbdb46",
      "name": "node-2"
    },
    "disk_class": "solid_state",
```

```
    "disk_type": "ssd",
    "size": 1902379008,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partition": false,
    "usable": 1,
    "layout_requirements": [
      {
        "raid_type": "raid4",
        "default": true,
        "aggregate_min_disks": 2,
        "raid_group": {
          "min": 2,
          "max": 14,
          "default": 8
        }
      }
    ]
  }
]
}
----

=== Retrieving a recommendation for an aggregate expansion

The following example shows the response with the recommended data
aggregate expansion based on what disks are present within the system.
The query does not modify the existing aggregate but returns how the
aggregate will look after the expansion. The recommendation will be
reflected in the attributes - 'space.block_storage.size' and
'block_storage.primary.disk_count'.
Recommended data aggregate expansion will be blocked while one or more
nodes in the cluster are simulating or implementing automatic aggregate
creation.

----

# The API:
/api/storage/aggregates/{uuid}?auto_provision_policy=expand

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-
babb-ef437d118314?auto_provision_policy=expand" -H "accept:
application/json"

# The response:
```

```
{
"uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
"name": "node_2_SSD_1",
"node": {
  "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
  "name": "node-2"
},
"space": {
  "block_storage": {
    "size": 1116180480
  }
},
"block_storage": {
  "primary": {
    "disk_count": 12,
    "disk_class": "solid_state",
    "raid_type": "raid_dp",
    "disk_type": "ssd",
    "raid_size": 24,
    "simulated_raid_groups": [
       {
         "name": "test/plex0/rg0",
         "parity_disk_count": 2,
         "data_disk_count": 10,
         "usable_size": 12309487,
         "is_partition": false
       }
     ]
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false
  }
}
}
----

=== Updating an aggregate in the cluster

The following example shows the workflow of adding disks to the aggregate.

Step 1: Check the current disk count on the aggregate.

----
```

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-
8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept:
application/json"

# The response:
{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "test1",
"block_storage": {
  "primary": {
    "disk_count": 6
  }
},
}
----

Step 2: Update the aggregate with the new disk count in
'block_storage.primary.disk_count'. The response to PATCH is a job unless
the request is invalid.

----

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-
4a9f-8f01-712f626c983c" -H "accept: application/hal+json" -d
"{\"block_storage\": {\"primary\": {\"disk_count\": 8}}}"

# The response:
{
"job": {
  "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
    }
  }
}
}
```

```
----

Step 3: Wait for the job to finish, then call GET to see the reflected
change.

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-
8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept:
application/json"

# The response:
{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "test1",
"block_storage": {
  "primary": {
    "disk_count": 8
  }
},
}
----

=== Adding a storage pool to an aggregate

The following example shows how to add cache capacity from an existing
storage pool to an aggregate.
Step 1: Update the aggregate with the new storage pool allocation unit in
'block_storage.hybrid_cache.storage_pools.allocation_units_count'.
Additionally, specify 'block_storage.hybrid_cache.storage_pools.name' or
'block_storage.hybrid_cache.storage_pools.uuid' to the storage pool.
On the first storage pool, 'block_storage.hybrid_cache.raid_type' can be
specified for the raidtype of the hybrid cache.
The response to PATCH is a job unless the request is invalid.

----

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-
4a9f-8f01-712f626c983c" -H "accept: application/json" -d
```

```
"{\"block_storage\": {\"hybrid_cache\": {\"raid_type\": \"raid_dp\",
\"storage_pools\": [{ \"allocation_units_count\": 2, \"storage_pool\": {
\"name\": \"sp1\"}}]}}}"

# The response:
{
"job": {
  "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
    }
  }
}
}
----

Step 2: Wait for the job to finish, then call GET to see the reflected
change.

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-
8f01-712f626c983c?fields=block_storage.hybrid_cache" -H "accept:
application/json"

# The response:
{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "test1",
"hybrid_cache": {
    "enabled": true,
    "disk_count": 3,
    "raid_size": 24,
    "raid_type": "raid_dp",
    "size": 880279552,
    "used": 73728,
    "storage_pools": [
        {
            "allocation_units_count": 2,
            "storage_pool": {
                "name": "sp1",
```

```
                "uuid": "eeef0b24-846b-11ec-8fcb-005056bb12c7"
            }
        }
    ]
}
}
----

=== Adding physical SSD cache capacity to an aggregate

The following example shows how to add physical SSD cache capacity to an
aggregate.
Step 1: Specify the number of disks to be added to cache in
'block_storage.hybrid_cache.disk_count'.
'block_storage.hybrid_cache.raid_type' can be specified for the RAID type
of the hybrid cache.
'block_storage.hybrid_cache.raid_size' can be specified for the RAID size
of the hybrid cache.
The response to PATCH is a job unless the request is invalid.

----

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/caa8a9f1-0219-
4eaf-bcad-e29c05042fe1" -H "accept: application/json" -d
'{"block_storage.hybrid_cache.disk_count":3,"block_storage.hybrid_cache.ra
id_type":"raid4"}'

# The response:
{
"job": {
  "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
    }
  }
}
}
----

Step 2: Wait for the job to finish, then call GET to see the reflected
change.
```

```
----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/caa8a9f1-0219-4eaf-
bcad-e29c05042fe1?fields=block_storage.hybrid_cache" -H "accept:
application/json"

# The response:
{
"uuid": "caa8a9f1-0219-4eaf-bcad-e29c05042fe1",
"name": "test1",
"hybrid_cache": {
    "enabled": true,
    "disk_count": 3,
    "raid_size": 24,
    "raid_type": "raid4",
    "size": 880279552,
    "used": 73728
}
}
----

=== Simulated addition of physical SSD cache capacity to an aggregate

The following example shows the response for a simulated addition of
physical SSD cache capacity to an aggregate based on the values of the
'block_storage.hybrid_cache.disk_count',
'block_storage.hybrid_cache.raid_type' and
'block_storage.hybrid_cache.raid_size' attributes passed in.
The query does not modify the existing aggregate but returns how the
aggregate will look after the expansion along with any associated
warnings.
Simulated addition of physical SSD cache capacity to an aggregate will be
blocked while one or more nodes in the cluster are simulating or
implementing automatic aggregate creation.
This will be reflected in the following attributes:

* block_storage.hybrid_cache.size - Total usable cache space in bytes, not
including WAFL reserve and aggregate snapshot reserve.
* block_storage.hybrid_cache.disk_count - Number of disks that can be
added to the aggregate's cache tier.

----
```

```
# The API:
/api/storage/aggregates/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/7eb630d1-0e55-
4cb6-8d90-957d6f4db54e?simulate=true" -H "accept: application/json" -d
'{"block_storage.hybrid_cache.disk_count":6,"block_storage.hybrid_cache.ra
id_type":"raid4","block_storage.hybrid_cache.raid_size":3}'

# The response:
{
 "warnings": [
   {
     "name": "test",
     "warning": {
       "code": 18316,
       "message": "Operation will lead to creation of new raid group"
     }
   }
 ],
 "num_records": 1,
 "records": [
   {
     "uuid": "7eb630d1-0e55-4cb6-8d90-957d6f4db54e",
     "name": "test",
     "node": {
       "uuid": "30d69eb5-f046-11ec-9bba-005056bba492",
       "name": "node-1",
       "_links": {
         "self": {
           "href": "/api/cluster/nodes/30d69eb5-f046-11ec-9bba-
005056bba492"
         }
       }
     },
     "space": {
       "block_storage": {
         "size": 833777664
       }
     },
     "block_storage": {
       "primary": {
         "disk_count": 3,
         "disk_class": "virtual",
         "raid_type": "raid_dp",
         "disk_type": "vm_disk"
```

```
        },
        "hybrid_cache": {
          "disk_class": "solid_state",
          "disk_type": "ssd",
          "enabled": false,
          "disk_count": 6,
          "raid_type": "raid4",
          "size": 6771179520,
          "simulated_raid_groups": [
            {
              "name": "/test/plex0/rg1",
              "existing_parity_disk_count": 0,
              "existing_data_disk_count": 0,
              "added_parity_disk_count": 1,
              "added_data_disk_count": 2,
              "usable_size": 1880883200,
              "is_partition": false
            },
            {
              "name": "/test/plex0/rg2",
              "existing_parity_disk_count": 0,
              "existing_data_disk_count": 0,
              "added_parity_disk_count": 1,
              "added_data_disk_count": 2,
              "usable_size": 1880883200,
              "is_partition": false
            }
          ]
        },
        "mirror": {
          "enabled": false
        }
      },
      "_links": {
        "self": {
          "href": "/api/storage/aggregates/7eb630d1-0e55-4cb6-8d90-
957d6f4db54e"
        }
      }
    }
  ]
  }
----

The following example shows the workflow to enable software encryption on
an aggregate.
```

```
Step 1: Check the current software encryption status of the aggregate.

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-
9590-5a402bffbe4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

# The response:
{
"uuid": "f3aafdc6-be35-4d93-9590-5a402bffbe4b",
"name": "aggr5",
"data_encryption": {
  "software_encryption_enabled": false
},
}
----

Step 2: Update the aggregate with the encryption status in
'data_encryption.software_encryption_enabled'. The response to PATCH is a
job unless the request is invalid.

----

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-
4d93-9590-5a402bffbe4b" -H "accept: application/hal+json" -d
'{"data_encryption": {"software_encryption_enabled": "true"}}'

# The response:
{
"job": {
  "uuid": "6b7ab28e-168d-11ea-8a50-0050568eca76",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6b7ab28e-168d-11ea-8a50-0050568eca76"
    }
  }
```

```
}
}
----

Step 3: Wait for the job to finish, then call GET to see the reflected
change.

----

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-
9590-5a402bffbe4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

# The response:
{
"uuid": "f3aafdc6-be35-4d93-9590-5a402bffbe4b",
"name": "aggr5",
"data_encryption": {
  "software_encryption_enabled": true
},
}
----




[[ID420df642d59b91b8e8c4374ee62150c4]]
= Delete an aggregate specified by the UUID

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/aggregates/{uuid}`#

*Introduced In:* 9.6

Deletes the aggregate specified by the UUID. This request starts a job and
returns a link to that job.

== Related ONTAP commands

* `storage aggregate delete`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Aggregate UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 460770
| The aggregate delete job failed to delete the aggregate.
```

```
| 460777
| Failed to get information on the delete job.

| 786435
| Internal Error. Failed to create a communication handle.

| 786451
| Failed to delete specified aggregate.

| 786468
| VLDB is offline.

| 786472
| Node that hosts the aggregate is offline.

| 786497
| Cannot delete an aggregate that has volumes.

| 786771
| Aggregate does not exist.

| 786867
| Specified aggregate resides on the remote cluster.

| 786897
| Specified aggregate cannot be deleted as it is a switched-over root
aggregate.

| 19726544
| DELETE on the aggregate endpoint is not supported on this version of
ONTAP.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID6f116ebe87cf714941540122a055d6a]]
= Retrieve an aggregate specified by the UUID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{uuid}`#

*Introduced In:* 9.6

Retrieves the aggregate specified by the UUID. The recommend query cannot
be used for this operation.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `metric.+*+`
* `space.block_storage.inactive_user_data`
* `space.block_storage.inactive_user_data_percent`
* `space.footprint`
* `is_spare_low`
* `statistics.+*+`

== Related ONTAP commands

* `storage aggregate show`


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Aggregate UUID


|show_spares
|boolean
|query
|False
a|If set to 'true', the spares object is returned instead of records to
show the spare groups in the cluster that are compatible with the
aggregate. The default setting is 'false'.

* Introduced in: 9.12


|flash_pool_eligible
|boolean
|query
|False
a|If set to 'true' along with show_spares, the spares object is restricted
to return spare groups that are compatible with flash pool creation or
addition for the specified aggregate. The default setting is 'false'.

* Introduced in: 9.12


|auto_provision_policy
|string
|query
|False
a|If set to expand, a query is run on the system for the recommended
optimal expansion layout of the aggregate.

* Introduced in: 9.8


|fields
|array[string]
|query
```

```
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|cloud_storage
|link:#cloud_storage[cloud_storage]
a|Configuration information for the cloud storage portion of the
aggregate. This is referred to as the capacity tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|
```

```
|dr_home_node
|link:#dr_home_node[dr_home_node]
a|Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.


|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|

|inode_attributes
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the aggregate.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
```

```
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snaplock_type
|string
a|SnapLock type.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
|string
a|Operational state of the aggregate.


|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the aggregate.


|uuid
|string
a|Aggregate UUID.


|volume-count
|integer
a|Number of volumes in the aggregate.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "string",
      "raid_size": 24,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string"
        }
      ],
      "size": 1612709888,
      "storage_pools": [
        {
          "storage_pool": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "storage_pool_1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "plex0"
```

```
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "disk_type": "string",
      "raid_size": 16,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string",
          "raid_type": "string"
        }
      ]
    },
    "storage_type": "string"
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 12:00:00 -0400",
  "dr_home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "inactive_data_reporting": {
    "start_time": "2019-12-12 12:00:00 -0400"
  },
  "inode_attributes": {
    "file_private_capacity": 31136,
    "file_public_capacity": 31136,
    "files_private_used": 502,
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "node1_aggr_1",
  "node": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snaplock_type": "string",
  "snapshot": {
    "files_total": 10,
    "files_used": 3,
    "max_files_available": 5,
    "max_files_used": 50
  },
  "space": {
    "block_storage": {
      "aggregate_metadata": 2655,
      "aggregate_metadata_percent": 8,
      "available": 10156560384,
      "data_compacted_count": 1990000,
      "data_compaction_space_saved": 1996000,
      "data_compaction_space_saved_percent": 27,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "inactive_user_data_percent": 0,
      "physical_used": 2461696,
      "physical_used_percent": 50,
      "size": 10156769280,
      "used": 2088960,
      "used_including_snapshot_reserve": 674685,
      "used_including_snapshot_reserve_percent": 35,
      "used_percent": 50,
      "volume_deduplication_shared_count": 1990000,
      "volume_deduplication_space_saved": 1996000,
      "volume_deduplication_space_saved_percent": 27,
      "volume_footprints_percent": 14
    },
    "cloud_storage": {
      "used": 402743264
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0,
      "wise_tsse_min_used_capacity_pct": 0
    },
```

```
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots_flexclones": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": 608896,
    "snapshot": {
      "available": 2000,
      "reserve_percent": 20,
      "total": 5000,
      "used": 3000,
      "used_percent": 45
    },
    "total_provisioned_space": 20971520
  },
  "state": "string",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "uuid": "string",
  "volume-count": 0
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 787092
| The target field cannot be specified for this operation.

| 7209049
| Cannot perform the operation because the aggregate is currently
expanding.

| 8586225
| Unexpected error encountered when retrieving metrics and statistics for
this aggregate.

| 19726382
| Another provisioning operation is in progress on this cluster. Wait a
few minutes, and try the operation again.

| 19726390
| Unable to provide a recommendation to expand the aggregate.

| 19726391
| Too many unassigned disks visible to the node that owns this aggregate.

| 19726392
| Layout of this aggregate is not a supported configuration.

| 19726393
| Failed to expand the aggregate. Aggregate expansion is not supported on
this system.

| 19726394
| Automatic aggregate expansion is not supported on systems with multiple
data aggregates.

| 19726395
| Automatic aggregate expansion is not supported when MetroCluster is not
configured

| 19726396
| Automatic aggregate expansion is not supported when the DR group is not
in a normal state
```

```
| 19726397
| Aggregates must contain disks with identical disk-types and disk-sizes.

| 19726402
| Internal error. Unable to determine the MetroCluster configuration
state.

| 19726538
| Cannot perform the operation because the aggregate is not in a healthy
state.

| 19726541
| Cannot perform the operation because the specified aggregate is a root
aggregate.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
```

```
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|usable_size
|integer
a|Usable size of each disk, in bytes.
```

```
|===


[#storage_pool_reference]
[.api-collapsible-fifth-title]
storage_pool_reference

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pools]
[.api-collapsible-fifth-title]
storage_pools

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allocation_units_count
|integer
a|Allocation count of storage pool.


|storage_pool
```

```
|link:#storage_pool_reference[storage_pool_reference]
a|Shared Storage Pool


|===


[#hybrid_cache]
[.api-collapsible-fifth-title]
hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made
up of either whole SSDs or storage pool SSDs.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|Number of disks used in the cache tier of the aggregate. Only provided
when hybrid_cache.enabled is 'true'.


|disk_type
|string
a|Type of disk being used by the aggregate's cache tier.


|enabled
|boolean
a|Specifies whether the aggregate uses HDDs with SSDs as a cache.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type for SSD cache of the aggregate. Only provided when
hybrid_cache.enabled is 'true'.
```

```
|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|size
|integer
a|Total usable space in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|storage_pools
|array[link:#storage_pools[storage_pools]]
a|List of storage pool properties and allocation_units_count for
aggregate.


|used
|integer
a|Space used in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|===


[#mirror]
[.api-collapsible-fifth-title]
mirror

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Aggregate is SyncMirror protected


|state
|string
a|

|===
```

```
[#plex_reference]
[.api-collapsible-fifth-title]
plex_reference

Plex


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===


[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|data_disk_count
```

```
|integer
a|Number of data disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|parity_disk_count
|integer
a|Number of parity disks in RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#primary]
[.api-collapsible-fifth-title]
primary

Configuration information for the primary storage portion of the
```

```
aggregate. This excludes the hybrid cache details.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|checksum_style
|string
a|The checksum style used by the aggregate.


|disk_class
|string
a|The class of disks being used by the aggregate.


|disk_count
|integer
a|Number of disks used in the aggregate. This includes parity disks, but
excludes disks in the hybrid cache.


|disk_type
|string
a|The type of disk being used by the aggregate.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|===
```

```
[#block_storage]
[.api-collapsible-fifth-title]
block_storage

Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|hybrid_cache
|link:#hybrid_cache[hybrid_cache]
a|Contains the configuration for the hybrid cache. The hybrid cache is
made up of either whole SSDs or storage pool SSDs.


|mirror
|link:#mirror[mirror]
a|

|plexes
|array[link:#plex_reference[plex_reference]]
a|Plex reference for each plex in the aggregate.


|primary
|link:#primary[primary]
a|Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


|storage_type
|string
a|Type of aggregate.


|uses_partitions
|boolean
a|If true, aggregate is using shared disks.
```

```
|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

Cloud store


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#cloud_storage_tier]
[.api-collapsible-fifth-title]
cloud_storage_tier

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cloud_store
|link:#cloud_store[cloud_store]
a|Cloud store


|used
```

```
|integer
a|Capacity used in bytes in the cloud store by this aggregate. This is a
cached value calculated every 5 minutes.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

Configuration information for the cloud storage portion of the aggregate.
This is referred to as the capacity tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|attach_eligible
|boolean
a|Specifies whether the aggregate is eligible for a cloud store to be
attached.


|migrate_threshold
|integer
a|Specifies the minimum percentage of performance tier free space that
must exist in order for migration of data from the capacity tier to
performance tier to be allowed. Only valid for PATCH operations.


|stores
|array[link:#cloud_storage_tier[cloud_storage_tier]]
a|Configuration information for each cloud storage portion of the
aggregate.


|tiering_fullness_threshold
|integer
a|The percentage of space in the performance tier that must be used before
data is tiered out to the cloud store. Only valid for PATCH operations.
```

```
|===


[#data_encryption]
[.api-collapsible-fifth-title]
data_encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drive_protection_enabled
|boolean
a|Specifies whether the aggregate uses self-encrypting drives with data
protection enabled.


|software_encryption_enabled
|boolean
a|Specifies whether NetApp aggregate encryption is enabled. All data in
the aggregate is encrypted.


|===


[#dr_home_node]
[.api-collapsible-fifth-title]
dr_home_node

Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
```

```
|string
a|

|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

Node where the aggregate resides after giveback. The value for this field
might differ from the value of the 'node' field during takeover.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#inactive_data_reporting]
[.api-collapsible-fifth-title]
inactive_data_reporting

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
```

```
a|Specifies whether or not inactive data reporting is enabled on the
aggregate.


|start_time
|string
a|Timestamp at which inactive data reporting was enabled on the aggregate.


|===


[#inode_attributes]
[.api-collapsible-fifth-title]
inode_attributes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|file_private_capacity
|integer
a|Number of files that can currently be stored on disk for system metadata
files. This number will dynamically increase as more system files are
created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|file_public_capacity
|integer
a|Number of files that can currently be stored on disk for user-visible
files.  This number will dynamically increase as more user-visible files
are created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_private_used
|integer
a|Number of system metadata files used. If the referenced file system is
```

```
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_total
|integer
a|Maximum number of user-visible files that this referenced file system
can currently hold. If the referenced file system is restricted or
offline, a value of 0 is returned.


|files_used
|integer
a|Number of user-visible files used in the referenced file system. If the
referenced file system is restricted or offline, a value of 0 is returned.


|max_files_available
|integer
a|The count of the maximum number of user-visible files currently
allowable on the referenced file system.


|max_files_possible
|integer
a|The largest value to which the maxfiles-available parameter can be
increased by reconfiguration, on the referenced file system.


|max_files_used
|integer
a|The number of user-visible files currently in use on the referenced file
system.


|used_percent
|integer
a|The percentage of disk space currently in use based on user-visible file
count on the referenced file system.


|version
|integer
a|The inofile-version of the aggregate. If the referenced file system is
```

```
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

The most recent sample of I/O metrics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
```

```
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node where the aggregate currently resides.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|files_total
|integer
a|Total files allowed in snapshots


|files_used
|integer
a|Total files created in snapshots


|max_files_available
|integer
a|Maximum files available for snapshots
```

```
|max_files_used
|integer
a|Files in use by snapshots


|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate_metadata
|integer
a|Space used by different metafiles and internal operations inside the
aggregate, in bytes.


|aggregate_metadata_percent
|integer
a|Aggregate metadata as a percentage.


|available
|integer
a|Space available in bytes.


|data_compacted_count
|integer
a|Amount of compacted data in bytes.


|data_compaction_space_saved
|integer
a|Space saved in bytes by compacting the data.


|data_compaction_space_saved_percent
|integer
```

```
a|Percentage saved by compacting the data.


|full_threshold_percent
|integer
a|The aggregate used percentage at which 'monitor.volume.full' EMS is
generated.


|inactive_user_data
|integer
a|The size that is physically used in the block storage and has a cold
temperature, in bytes. This property is only supported if the aggregate is
either attached to a cloud store or can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data or **.


|inactive_user_data_percent
|integer
a|The percentage of inactive user data in the block storage. This property
is only supported if the aggregate is either attached to a cloud store or
can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data_percent or **.


|physical_used
|integer
a|Total physical used size of an aggregate in bytes.


|physical_used_percent
|integer
a|Physical used percentage.


|size
|integer
a|Total usable space in bytes, not including WAFL reserve and aggregate
snapshot reserve.
```

```
|used
|integer
a|Space used or reserved in bytes. Includes volume guarantees and
aggregate metadata.


|used_including_snapshot_reserve
|integer
a|Total used including the snapshot reserve, in bytes.


|used_including_snapshot_reserve_percent
|integer
a|Total used including the Snapshot reserve as a percentage.


|used_percent
|integer
a|Aggregate used percentage.


|volume_deduplication_shared_count
|integer
a|Amount of shared bytes counted by storage efficiency.


|volume_deduplication_space_saved
|integer
a|Amount of space saved in bytes by storage efficiency.


|volume_deduplication_space_saved_percent
|integer
a|Percentage of space saved by storage efficiency.


|volume_footprints_percent
|integer
a|A summation of volume footprints inside the aggregate, as a percentage.
A volume's footprint is the amount of space being used for the volume in
the aggregate.


|===
```

```
[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|used
|integer
a|Used space in bytes in the cloud store. Only applicable for aggregates
with a cloud store tier.


|===


[#efficiency]
[.api-collapsible-fifth-title]
efficiency

Storage efficiency.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_adaptive_compression_savings
|boolean
a|Indicates whether or not aggregate has auto adaptive compression
savings.


|cross_volume_background_dedupe
|boolean
a|Indicates whether or not cross volume background deduplication is
enabled.


|cross_volume_dedupe_savings
|boolean
```

```
a|Indicates whether or not aggregate has cross volume deduplication
savings.


|cross_volume_inline_dedupe
|boolean
a|Indicates whether or not cross volume inline deduplication is enabled.


|enable_workload_informed_tsse
|boolean
a|Indicates whether Workload Informed TSSE is enabled on the system.


|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|wise_tsse_min_used_capacity_pct
|integer
a|Minimum amount of used data in aggregate required to trigger cold
compression on TSSE volumes.


|===


[#efficiency_without_snapshots]
[.api-collapsible-fifth-title]
efficiency_without_snapshots

Storage efficiency that does not include the savings provided by
snapshots.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots_flexclones]
[.api-collapsible-fifth-title]
efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by snapshots
and Flexclone volumes.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)
```

```
|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|Available space for snapshots in bytes


|reserve_percent
|integer
a|Percentage of space reserved for snapshots


|total
|integer
a|Total space for snapshots in bytes


|used
|integer
a|Space used by snapshots in bytes


|used_percent
|integer
a|Percentage of disk space used by snapshots


|===


[#space]
```

```
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_storage
|link:#block_storage[block_storage]
a|

|cloud_storage
|link:#cloud_storage[cloud_storage]
a|

|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and Flexclone volumes.


|footprint
|integer
a|A summation of volume footprints (including volume guarantees), in
bytes. This includes all of the volume footprints in the block_storage
tier and the cloud_storage tier.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|snapshot
```

```
|link:#snapshot[snapshot]
a|

|total_provisioned_space
|integer
a|A summation of volume size, in bytes. This includes all of the volumes
inside the aggregate.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either total_provisioned_space or **.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

The real time I/O statistics for the aggregate.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


|timestamp
|string
a|The timestamp of the performance data.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
```

```
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID24603294142a7eb02ca98ab7c1bba64f]]
= Update an aggregate specified by the UUID

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/aggregates/{uuid}`#

*Introduced In:* 9.6

Updates the aggregate specified by the UUID with the properties in the
body. This request starts a job and returns a link to that job.

== Related ONTAP commands

* `storage aggregate add-disks`
* `storage aggregate mirror`
* `storage aggregate modify`
* `storage aggregate relocation start`
* `storage aggregate rename`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|Aggregate UUID


|auto_provision_policy
|string
|query
|False
a|If set to expand, the PATCH operation runs the recommended expansion of
the aggregate.

* Introduced in: 9.8


|simulate
|boolean
|query
|False
a|If set to true, the PATCH operation runs a simulated aggregate expansion
with the provided input disk count and returns the proposed size of the
new aggregate along with any associated warnings.

* Introduced in: 9.8


|disk_size
|integer
|query
|False
a|If set, PATCH only selects disks of the specified size.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.
```

```
* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|cloud_storage
|link:#cloud_storage[cloud_storage]
a|Configuration information for the cloud storage portion of the
aggregate. This is referred to as the capacity tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|

|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
```

```
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|

|inode_attributes
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
|string
a|Operational state of the aggregate.
```

```
|uuid
|string
a|Aggregate UUID.


|volume-count
|integer
a|Number of volumes in the aggregate.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "string",
      "raid_size": 24,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string"
        }
      ],
      "size": 1612709888,
      "storage_pools": [
        {
          "storage_pool": {
            "name": "storage_pool_1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "used": 26501122
    },
```

```
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "name": "plex0"
      }
    ],
    "primary": {
      "disk_count": 8,
      "disk_type": "string",
      "raid_size": 16,
      "raid_type": "string",
      "simulated_raid_groups": [
        {
          "name": "string",
          "raid_type": "string"
        }
      ]
    },
    "storage_type": "string"
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 12:00:00 -0400",
  "home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "inactive_data_reporting": {
    "start_time": "2019-12-12 12:00:00 -0400"
  },
  "inode_attributes": {
    "file_private_capacity": 31136,
    "file_public_capacity": 31136,
    "files_private_used": 502,
```

```
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "name": "node1_aggr_1",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapshot": {
    "files_total": 10,
    "files_used": 3,
    "max_files_available": 5,
    "max_files_used": 50
  },
  "space": {
    "block_storage": {
      "aggregate_metadata": 2655,
      "aggregate_metadata_percent": 8,
      "available": 10156560384,
      "data_compacted_count": 1990000,
      "data_compaction_space_saved": 1996000,
      "data_compaction_space_saved_percent": 27,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "inactive_user_data_percent": 0,
      "physical_used": 2461696,
      "physical_used_percent": 50,
      "size": 10156769280,
      "used": 2088960,
      "used_including_snapshot_reserve": 674685,
      "used_including_snapshot_reserve_percent": 35,
      "used_percent": 50,
      "volume_deduplication_shared_count": 1990000,
      "volume_deduplication_space_saved": 1996000,
      "volume_deduplication_space_saved_percent": 27,
      "volume_footprints_percent": 14
    },
    "cloud_storage": {
      "used": 402743264
    },
```

```
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0,
      "wise_tsse_min_used_capacity_pct": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots_flexclones": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": 608896,
    "snapshot": {
      "available": 2000,
      "reserve_percent": 20,
      "total": 5000,
      "used": 3000,
      "used_percent": 45
    },
    "total_provisioned_space": 20971520
  },
  "state": "string",
  "uuid": "string",
  "volume-count": 0
}
====

== Response
```


Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#aggregate[aggregate]]
a|

|warnings
|array[link:#aggregate_warning[aggregate_warning]]
a|List of validation warnings and remediation advice for the aggregate
simulate behavior.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  },
  "num_records": 1,
  "records": [
    {
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "block_storage": {
        "hybrid_cache": {
          "disk_count": 6,
          "disk_type": "string",
          "raid_size": 24,
          "raid_type": "string",
          "simulated_raid_groups": [
            {
              "name": "string"
            }
          ],
          "size": 1612709888,
          "storage_pools": [
```

```
            {
              "storage_pool": {
                "name": "storage_pool_1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
              }
            }
          ],
          "used": 26501122
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
        "plexes": [
          {
            "name": "plex0"
          }
        ],
        "primary": {
          "disk_count": 8,
          "disk_type": "string",
          "raid_size": 16,
          "raid_type": "string",
          "simulated_raid_groups": [
            {
              "name": "string",
              "raid_type": "string"
            }
          ]
        },
        "storage_type": "string"
      },
      "cloud_storage": {
        "stores": [
          {
            "cloud_store": {
              "name": "store1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "used": 0
          }
        ]
      },
      "create_time": "2018-01-01 12:00:00 -0400",
      "home_node": {
        "name": "node1",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "inactive_data_reporting": {
        "start_time": "2019-12-12 12:00:00 -0400"
      },
      "inode_attributes": {
        "file_private_capacity": 31136,
        "file_public_capacity": 31136,
        "files_private_used": 502,
        "files_total": 31136,
        "files_used": 97,
        "max_files_available": 31136,
        "max_files_possible": 2844525,
        "max_files_used": 97,
        "used_percent": 5,
        "version": 4
      },
      "is_spare_low": "",
      "name": "node1_aggr_1",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snapshot": {
        "files_total": 10,
        "files_used": 3,
        "max_files_available": 5,
        "max_files_used": 50
      },
      "space": {
        "block_storage": {
          "aggregate_metadata": 2655,
          "aggregate_metadata_percent": 8,
          "available": 10156560384,
          "data_compacted_count": 1990000,
          "data_compaction_space_saved": 1996000,
          "data_compaction_space_saved_percent": 27,
          "full_threshold_percent": 0,
          "inactive_user_data": 304448,
          "inactive_user_data_percent": 0,
          "physical_used": 2461696,
          "physical_used_percent": 50,
          "size": 10156769280,
          "used": 2088960,
          "used_including_snapshot_reserve": 674685,
          "used_including_snapshot_reserve_percent": 35,
```

```
          "used_percent": 50,
          "volume_deduplication_shared_count": 1990000,
          "volume_deduplication_space_saved": 1996000,
          "volume_deduplication_space_saved_percent": 27,
          "volume_footprints_percent": 14
        },
        "cloud_storage": {
          "used": 402743264
        },
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0,
          "wise_tsse_min_used_capacity_pct": 0
        },
        "efficiency_without_snapshots": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots_flexclones": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "footprint": 608896,
        "snapshot": {
          "available": 2000,
          "reserve_percent": 20,
          "total": 5000,
          "used": 3000,
          "used_percent": 45
        },
        "total_provisioned_space": 20971520
      },
      "state": "string",
      "uuid": "string",
      "volume-count": 0
    }
  ],
  "warnings": [
    {
      "action": {
        "arguments": [
          "string"
        ],
```

```
        "message": "string"
      },
      "name": "string",
      "warning": {
        "arguments": [
          "string"
        ],
        "message": "string"
      }
    }
  ]
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262247
| The value is invalid for the field.

| 460777
| Failed to get information on the job.

| 786434
| Cannot connect to node where the aggregate resides.

| 786435
| Internal Error. Failed to create a communication handle.

| 786439
| An aggregate already uses the specified name.

| 786447
| Failed to modify the aggregate.
```

```
| 786456
| Failed to add disks to the aggregate.

| 786458
| Failed to rename aggregate.

| 786465
| Failed to take aggregate offline because of the reason outlined in the
message.

| 786467
| Failed to take aggregate online because of the reason outlined in the
message.

| 786468
| VLDB is offline.

| 786472
| Node that hosts the aggregate is offline.

| 786479
| Cannot find node ID for the node.

| 786491
| Not enough spares on the node.

| 786730
| Internal Error

| 786771
| Aggregate does not exist.

| 786787
| Aggregate is not online.

| 786808
| Aggregate mirror failed.

| 786867
| Specified aggregate resides on the remote cluster.

| 786911
| Not every node in the cluster has the Data ONTAP version required for
the feature.

| 786923
| This operation is not allowed during the pre-commit phase of a 7-mode to
```

```
clustered Data ONTAP transition.

| 786924
| Internal error for an aggregate that is in the pre-commit phase of a 7-
mode to clustered Data ONTAP transition.

| 786955
| Modifying raidtype to raid_tec requires a minimum of six disks in the
RAID Group.

| 786956
| Modifying raidtype to raid_dp requires a minimum of four disks in the
RAID Group.

| 786965
| Spare Selection in userspace failed.

| 787046
| Mirroring of a FabricPool is not allowed.

| 787092
| The target field cannot be specified for this operation.

| 787144
| Aggregate is not a FabricPool.

| 787156
| Modifying the attributes of mirror object store is not allowed.

| 787169
| Only one field can be modified per operation.

| 787170
| Failed to patch the "block_storage.primary.disk_count" because the disk
count specified is smaller than existing disk count.

| 787172
| This query is only allowed during the modification of the specified
field.

| 787178
| Unmirroring an aggregate with a PATCH operation is not supported.

| 787187
| Internal error. Failed to check if the aggregate is a FabricPool.

| 787266
```

```
| Invalid aggregate state. This state is not supported for a PATCH
operation.

| 787273
| Allocation unit count is not valid.

| 787274
| Raidtype is not valid.

| 787275
| Patch request with multiple records is not valid.

| 787276
| Storage pool name and uuid do not match.

| 787277
| Storage pool name and uuid are empty.

| 787278
| Incorrect storage pool name specified.

| 787279
| Incorrect storage pool uuid specified.

| 787280
| Allocation unit count specified is smaller than the existing allocation
unit count.

| 787281
| Cannot modify RAID type of aggregate hybrid cache tier.

| 787282
| RAID group must be specified on a disk addition to an aggregate with
mixed RAID types.

| 787283
| RAID group must be specified on a disk addition to a Flash Pool
aggregate.

| 787284
| The specified RAID group uses capacity from one or more storage pools.

| 787287
| Cannot add physical SSD cache because aggregate uses cache capacity from
a storage pool.

| 787288
```

```
| Cannot add storage pool units because aggregate uses physical SSD cache.

| 787289
| Incorrect raid_group specified during first time addition of physical
SSD cache to an aggregate.

| 787291
| Cannot specify RAID group which is located on the aggregate primary
tier.

| 787293
| Cannot specify RAID group which is located on the aggregate cache tier.

| 787294
| This query is only allowed during the modification of the specific
fields.

| 787295
| The storage pool allocation units count is required.

| 1258699
| Cannot use all the disks specified for the requested operation.

| 1263500
| Operation will lead to creation of new raid group.

| 1263501
| Operation will exceed half of the maximum volume sizes allowed on the
node.

| 1263502
| One spare data partition from at least one of the chosen root-data1-
data2 disks will not be used.

| 1263503
| Operation will lead to downsizing of one or more disks.

| 1263504
| Operation will lead to a spares low condition.

| 1263598
| One or more selected disks will be partitioned.

| 1263624
| Operation will lead to a no sparecore condition.

| 2425736
```

```
| No matching node found for the UUID provided.

| 7208962
| Aggregate in an inconsistent state.

| 7208993
| Failed to offline as the volume is being used.

| 7209033
| CIFS open files prevent operation.

| 7209049
| Cannot perform the operation because the aggregate is currently
expanding.

| 7209075
| Cannot perform the operation because the volume size limit for this
system type would be exceeded.

| 7209090
| Inconsistent state.

| 7209183
| Volume is a partial volume.

| 7209229
| This version of ONTAP does not recognize the filesystem. It is probably
from a later version of the software and is being left offline.

| 7209246
| The specified operation could not be completed as the volume is
currently busy.

| 7209247
| The volume was not found.

| 7209263
| Container has failed.

| 7209271
| wafliron is currently active.

| 7209275
| Container was created in an unclustered ONTAP deployment.

| 7209463
| Nvfile replay pending.
```

```
| 7209966
| Another online request is already in progress for aggregate. The
previous online request is waiting on a response from the licensing
manager.

| 11206666
| Storage pool is not healthy.

| 11210659
| Aggregate is not online.

| 11210662
| Adding capacity from storage pool to a mirrored aggregate is not
supported.

| 11210667
| Storage pool does not have enough spare allocation units.

| 11210670
| Cannot add capacity from storage pool to aggregate, because currently
allocated capacity to the aggregate does not span across all disks
belonging to the storage pool.

| 11210672
| Cannot grow aggregate as no capacity is allocated to it from storage
pool.

| 11210673
| Mixing of physical SSDs and capacity from a storage pool is not allowed
in same aggregate.

| 11210675
| Capacity in storage pool belongs to different fault isolation domain
than aggregate.

| 11210678
| Storage pool does not have enough disks to create RAID groups of same
raid type as that of already allocated cache tier.

| 11210679
| Storage pool does not have enough disks to create RAID groups of type
RAID-DP.

| 11210680
| Storage pool does not have enough disks to create RAID groups of same
raid type as that of the aggregate.
```

```
| 11210685
| Storage pool does not have enough disks to create RAID groups of type
RAID-TEC.

| 11210688
| Capacity from storage pool cannot be added to an SSD aggregate and a
Flash Pool.

| 13108106
| Cannot run aggregate relocation because volume expand is in progress.

| 19726347
| There are a number of unassigned disks visible to the node that owns
this aggregate.

| 19726382
| Another provisioning operation is in progress on this cluster. Wait a
few minutes, and try the operation again.

| 19726390
| Unable to automatically expand this aggregate.

| 19726391
| Too many unassigned disks visible to the node that owns this aggregate.

| 19726392
| Layout of this aggregate is not a supported configuration.

| 19726393
| Failed to expand the aggregate. Aggregate expansion is not supported on
this system.

| 19726394
| Automatic aggregate expansion is not supported on systems with multiple
data aggregates.

| 19726395
| Automatic aggregate expansion is not supported when MetroCluster is not
configured.

| 19726396
| Automatic aggregate expansion is not supported when the DR group is not
in a normal state.

| 19726397
| Aggregates must contain disks with identical disk-types and disk-sizes.
```

```
| 19726402
| Internal error. Unable to determine the MetroCluster configuration
state.

| 19726538
| Cannot perform the operation because the aggregate is not in a healthy
state.

| 19726541
| Cannot perform the operation because the specified aggregate is a root
aggregate.

| 19726544
| PATCH on the aggregate endpoint is not supported on this version of
ONTAP.

| 26542083
| Destination node is at higher Data ONTAP version than source node.

| 26542084
| Source node is at higher Data ONTAP version than destination node.

| 26542097
| Unable to get D-blade ID of destination.

| 26542101
| Unable to contact the source node.

| 26542102
| Unable to contact the destination node.

| 26542120
| An SVM migrate operation is in progress. When the migrate operation
completes, try the operation again.

| 26542121
| A MetroCluster disaster recovery operation is in progress. When the
recovery operation completes, try the operation again.

| 196608334
| Failed to modify the aggregate because it contains NAE volumes.

| 196608335
| Failed to modify the aggregate because it contains non-encrypted
volumes.
|===
```

```
Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.
```

```
|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#storage_pool_reference]
[.api-collapsible-fifth-title]
storage_pool_reference

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|

|uuid
|string
a|

|===


[#storage_pools]
[.api-collapsible-fifth-title]
storage_pools

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allocation_units_count
|integer
a|Allocation count of storage pool.


|storage_pool
|link:#storage_pool_reference[storage_pool_reference]
a|Shared Storage Pool


|===


[#hybrid_cache]
[.api-collapsible-fifth-title]
hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made
up of either whole SSDs or storage pool SSDs.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|disk_count
|integer
a|Number of disks used in the cache tier of the aggregate. Only provided
when hybrid_cache.enabled is 'true'.


|disk_type
|string
a|Type of disk being used by the aggregate's cache tier.


|enabled
|boolean
a|Specifies whether the aggregate uses HDDs with SSDs as a cache.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.


|raid_type
|string
a|RAID type for SSD cache of the aggregate. Only provided when
hybrid_cache.enabled is 'true'.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|size
|integer
a|Total usable space in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|storage_pools
|array[link:#storage_pools[storage_pools]]
a|List of storage pool properties and allocation_units_count for
aggregate.


|used
|integer
```

```
a|Space used in bytes of SSD cache. Only provided when
hybrid_cache.enabled is 'true'.


|===


[#mirror]
[.api-collapsible-fifth-title]
mirror

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Aggregate is SyncMirror protected


|state
|string
a|

|===


[#plex_reference]
[.api-collapsible-fifth-title]
plex_reference

Plex


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===
```

```
[#simulated_raid_groups]
[.api-collapsible-fifth-title]
simulated_raid_groups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|added_data_disk_count
|integer
a|Number of added data disks in RAID group.


|added_parity_disk_count
|integer
a|Number of added parity disks in RAID group.


|data_disk_count
|integer
a|Number of data disks in RAID group.


|existing_data_disk_count
|integer
a|Number of existing data disks in the RAID group.


|existing_parity_disk_count
|integer
a|Number of existing parity disks in the RAID group.


|is_partition
|boolean
a|Indicates whether the disk is partitioned (true) or whole (false).


|name
|string
a|Name of the raid group.
```

```
|parity_disk_count
|integer
a|Number of parity disks in RAID group.


|raid_type
|string
a|RAID type of the aggregate.


|usable_size
|integer
a|Usable size of each disk, in bytes.


|===


[#primary]
[.api-collapsible-fifth-title]
primary

Configuration information for the primary storage portion of the
aggregate. This excludes the hybrid cache details.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|Number of disks used in the aggregate. This includes parity disks, but
excludes disks in the hybrid cache.


|disk_type
|string
a|The type of disk being used by the aggregate.


|raid_size
|integer
a|Option to specify the maximum number of disks that can be included in a
RAID group.
```

```
|raid_type
|string
a|RAID type of the aggregate.


|simulated_raid_groups
|array[link:#simulated_raid_groups[simulated_raid_groups]]
a|

|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|hybrid_cache
|link:#hybrid_cache[hybrid_cache]
a|Contains the configuration for the hybrid cache. The hybrid cache is
made up of either whole SSDs or storage pool SSDs.


|mirror
|link:#mirror[mirror]
a|

|plexes
|array[link:#plex_reference[plex_reference]]
a|Plex reference for each plex in the aggregate.


|primary
|link:#primary[primary]
a|Configuration information for the primary storage portion of the
```

```
aggregate. This excludes the hybrid cache details.


|storage_type
|string
a|Type of aggregate.


|uses_partitions
|boolean
a|If true, aggregate is using shared disks.


|===


[#cloud_store]
[.api-collapsible-fifth-title]
cloud_store

Cloud store


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#cloud_storage_tier]
[.api-collapsible-fifth-title]
cloud_storage_tier

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|cloud_store
|link:#cloud_store[cloud_store]
a|Cloud store


|used
|integer
a|Capacity used in bytes in the cloud store by this aggregate. This is a
cached value calculated every 5 minutes.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

Configuration information for the cloud storage portion of the aggregate.
This is referred to as the capacity tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|attach_eligible
|boolean
a|Specifies whether the aggregate is eligible for a cloud store to be
attached.


|migrate_threshold
|integer
a|Specifies the minimum percentage of performance tier free space that
must exist in order for migration of data from the capacity tier to
performance tier to be allowed. Only valid for PATCH operations.


|stores
|array[link:#cloud_storage_tier[cloud_storage_tier]]
a|Configuration information for each cloud storage portion of the
```

```
aggregate.


|tiering_fullness_threshold
|integer
a|The percentage of space in the performance tier that must be used before
data is tiered out to the cloud store. Only valid for PATCH operations.


|===


[#data_encryption]
[.api-collapsible-fifth-title]
data_encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drive_protection_enabled
|boolean
a|Specifies whether the aggregate uses self-encrypting drives with data
protection enabled.


|software_encryption_enabled
|boolean
a|Specifies whether NetApp aggregate encryption is enabled. All data in
the aggregate is encrypted.


|===


[#dr_home_node]
[.api-collapsible-fifth-title]
dr_home_node

Node where the aggregate resides after disaster recovery. The value for
this field might differ from the 'node' field during switchover.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

Node where the aggregate resides after giveback. The value for this field
might differ from the value of the 'node' field during takeover.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#inactive_data_reporting]
[.api-collapsible-fifth-title]
inactive_data_reporting

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|enabled
|boolean
a|Specifies whether or not inactive data reporting is enabled on the
aggregate.


|start_time
|string
a|Timestamp at which inactive data reporting was enabled on the aggregate.


|===


[#inode_attributes]
[.api-collapsible-fifth-title]
inode_attributes

[cols=3*,options=header]
|===
|Name
|Type
|Description

|file_private_capacity
|integer
a|Number of files that can currently be stored on disk for system metadata
files. This number will dynamically increase as more system files are
created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|file_public_capacity
|integer
a|Number of files that can currently be stored on disk for user-visible
files.  This number will dynamically increase as more user-visible files
are created.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.
```

```
|files_private_used
|integer
a|Number of system metadata files used. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|files_total
|integer
a|Maximum number of user-visible files that this referenced file system
can currently hold. If the referenced file system is restricted or
offline, a value of 0 is returned.


|files_used
|integer
a|Number of user-visible files used in the referenced file system. If the
referenced file system is restricted or offline, a value of 0 is returned.


|max_files_available
|integer
a|The count of the maximum number of user-visible files currently
allowable on the referenced file system.


|max_files_possible
|integer
a|The largest value to which the maxfiles-available parameter can be
increased by reconfiguration, on the referenced file system.


|max_files_used
|integer
a|The number of user-visible files currently in use on the referenced file
system.


|used_percent
|integer
a|The percentage of disk space currently in use based on user-visible file
count on the referenced file system.
```

```
|version
|integer
a|The inofile-version of the aggregate. If the referenced file system is
restricted or offline, a value of 0 is returned.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.
```

```
|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

The most recent sample of I/O metrics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
```

```
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#node]
[.api-collapsible-fifth-title]
node
```

```
Node where the aggregate currently resides.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|files_total
|integer
a|Total files allowed in snapshots


|files_used
|integer
a|Total files created in snapshots


|max_files_available
|integer
a|Maximum files available for snapshots


|max_files_used
|integer
```

```
a|Files in use by snapshots


|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate_metadata
|integer
a|Space used by different metafiles and internal operations inside the
aggregate, in bytes.


|aggregate_metadata_percent
|integer
a|Aggregate metadata as a percentage.


|available
|integer
a|Space available in bytes.


|data_compacted_count
|integer
a|Amount of compacted data in bytes.


|data_compaction_space_saved
|integer
a|Space saved in bytes by compacting the data.


|data_compaction_space_saved_percent
|integer
a|Percentage saved by compacting the data.
```

```
|full_threshold_percent
|integer
a|The aggregate used percentage at which 'monitor.volume.full' EMS is
generated.


|inactive_user_data
|integer
a|The size that is physically used in the block storage and has a cold
temperature, in bytes. This property is only supported if the aggregate is
either attached to a cloud store or can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data or **.


|inactive_user_data_percent
|integer
a|The percentage of inactive user data in the block storage. This property
is only supported if the aggregate is either attached to a cloud store or
can be attached to a cloud store.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either
block_storage.inactive_user_data_percent or **.


|physical_used
|integer
a|Total physical used size of an aggregate in bytes.


|physical_used_percent
|integer
a|Physical used percentage.


|size
|integer
a|Total usable space in bytes, not including WAFL reserve and aggregate
snapshot reserve.


|used
```

```
|integer
a|Space used or reserved in bytes. Includes volume guarantees and
aggregate metadata.


|used_including_snapshot_reserve
|integer
a|Total used including the snapshot reserve, in bytes.


|used_including_snapshot_reserve_percent
|integer
a|Total used including the Snapshot reserve as a percentage.


|used_percent
|integer
a|Aggregate used percentage.


|volume_deduplication_shared_count
|integer
a|Amount of shared bytes counted by storage efficiency.


|volume_deduplication_space_saved
|integer
a|Amount of space saved in bytes by storage efficiency.


|volume_deduplication_space_saved_percent
|integer
a|Percentage of space saved by storage efficiency.


|volume_footprints_percent
|integer
a|A summation of volume footprints inside the aggregate, as a percentage.
A volume's footprint is the amount of space being used for the volume in
the aggregate.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
```

```
cloud_storage

[cols=3*,options=header]
|===
|Name
|Type
|Description

|used
|integer
a|Used space in bytes in the cloud store. Only applicable for aggregates
with a cloud store tier.


|===


[#efficiency]
[.api-collapsible-fifth-title]
efficiency

Storage efficiency.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_adaptive_compression_savings
|boolean
a|Indicates whether or not aggregate has auto adaptive compression
savings.


|cross_volume_background_dedupe
|boolean
a|Indicates whether or not cross volume background deduplication is
enabled.


|cross_volume_dedupe_savings
|boolean
a|Indicates whether or not aggregate has cross volume deduplication
savings.
```

```
|cross_volume_inline_dedupe
|boolean
a|Indicates whether or not cross volume inline deduplication is enabled.


|enable_workload_informed_tsse
|boolean
a|Indicates whether Workload Informed TSSE is enabled on the system.


|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|wise_tsse_min_used_capacity_pct
|integer
a|Minimum amount of used data in aggregate required to trigger cold
compression on TSSE volumes.


|===


[#efficiency_without_snapshots]
[.api-collapsible-fifth-title]
efficiency_without_snapshots

Storage efficiency that does not include the savings provided by
snapshots.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots_flexclones]
[.api-collapsible-fifth-title]
efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by snapshots
and Flexclone volumes.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)
```

```
|===


[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|Available space for snapshots in bytes


|reserve_percent
|integer
a|Percentage of space reserved for snapshots


|total
|integer
a|Total space for snapshots in bytes


|used
|integer
a|Space used by snapshots in bytes


|used_percent
|integer
a|Percentage of disk space used by snapshots


|===


[#space]
[.api-collapsible-fifth-title]
space
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_storage
|link:#block_storage[block_storage]
a|

|cloud_storage
|link:#cloud_storage[cloud_storage]
a|

|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and Flexclone volumes.


|footprint
|integer
a|A summation of volume footprints (including volume guarantees), in
bytes. This includes all of the volume footprints in the block_storage
tier and the cloud_storage tier.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|snapshot
|link:#snapshot[snapshot]
a|
```

```
|total_provisioned_space
|integer
a|A summation of volume size, in bytes. This includes all of the volumes
inside the aggregate.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either total_provisioned_space or **.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.
```

```
|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

```
Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

The real time I/O statistics for the aggregate.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#aggregate]
```

```
[.api-collapsible-fifth-title]
aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
aggregate. When a cloud store is also used by this aggregate, this is
referred to as the performance tier.


|cloud_storage
|link:#cloud_storage[cloud_storage]
a|Configuration information for the cloud storage portion of the
aggregate. This is referred to as the capacity tier.


|create_time
|string
a|Timestamp of aggregate creation.


|data_encryption
|link:#data_encryption[data_encryption]
a|

|home_node
|link:#home_node[home_node]
a|Node where the aggregate resides after giveback. The value for this
field might differ from the value of the 'node' field during takeover.


|inactive_data_reporting
|link:#inactive_data_reporting[inactive_data_reporting]
a|
```

```
|inode_attributes
|link:#inode_attributes[inode_attributes]
a|

|is_spare_low
|boolean
a|Specifies whether the aggregate is in a spares low condition on any of
the RAID groups.
This is an advanced property; there is an added computational cost to
retrieving its value. The field is not populated for either a collection
GET or an instance GET unless it is explicitly requested using the
_fields_ query parameter containing either footprint or **.


|name
|string
a|Aggregate name.


|node
|link:#node[node]
a|Node where the aggregate currently resides.


|sidl_enabled
|boolean
a|Specifies whether or not SIDL is enabled on the aggregate.


|snapshot
|link:#snapshot[snapshot]
a|

|space
|link:#space[space]
a|

|state
|string
a|Operational state of the aggregate.


|uuid
|string
a|Aggregate UUID.
```

```
|volume-count
|integer
a|Number of volumes in the aggregate.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#action]
[.api-collapsible-fifth-title]
action

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the specified action message.


|code
|integer
a|Corrective action code of the specified action.
```

```
|message
|string
a|Specifies the corrective action to be taken to resolve the issue.


|===


[#warning]
[.api-collapsible-fifth-title]
warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[string]
a|Arguments present in the warning message encountered.


|code
|integer
a|Warning code of the warning encountered.


|message
|string
a|Details of the warning encountered by the aggregate simulate query.


|===


[#aggregate_warning]
[.api-collapsible-fifth-title]
aggregate_warning

[cols=3*,options=header]
|===
|Name
|Type
|Description

|action
```

```
|link:#action[action]
a|

|name
|string
a|Name of the entity that returns the warning.


|warning
|link:#warning[warning]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage aggregate metrics

:leveloffset: +1


[[ID6a35b23f88f60a625bed7f2764bad817]]
= Retrieve historical performance metrics for an aggregate

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/aggregates/{uuid}/metrics`#

*Introduced In:* 9.7

Retrieves historical performance metrics for an aggregate.
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|status
|string
|query
|False
a|Filter by status


|duration
|string
|query
|False
a|Filter by duration


|latency.other
|integer
|query
|False
a|Filter by latency.other


|latency.total
|integer
|query
|False
a|Filter by latency.total


|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read


|iops.other
|integer
|query
|False
a|Filter by iops.other


|iops.total
|integer
|query
|False
a|Filter by iops.total


|iops.write
|integer
|query
|False
a|Filter by iops.write


|iops.read
|integer
|query
|False
a|Filter by iops.read


|timestamp
|string
|query
|False
a|Filter by timestamp


|throughput.other
|integer
|query
|False
a|Filter by throughput.other
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total


|throughput.write
|integer
|query
|False
a|Filter by throughput.write


|throughput.read
|integer
|query
|False
a|Filter by throughput.read


|uuid
|string
|path
|True
a|Unique identifier of the aggregate.


|interval
|string
|query
|False
a|The time range for the data. Values can be 1h, 1d, 1w, 1m, or 1y.
The period for each time range is as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]


|return_timeout
```

```
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
```

```
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 8586225
| Encountered unexpected error in retrieving metrics for the requested
aggregate.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#iops]
[.api-collapsible-fifth-title]
iops
```

```
The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage bridges
```

```
:leveloffset: +1


[[IDb98e60b9045bf2f7db37de5d34295435]]
= Manage storage bridges



== Retrieving storage bridge information

The storage bridge GET API retrieves all of the bridges in the cluster.

'''

== Examples

=== 1) Retrieves a list of bridges from the cluster

==== The following example shows the response with a list of bridges from
the cluster:

'''

----

# The API:
/api/storage/bridges

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/bridges" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "wwn": "2000001086a18100",
    "name": "ATTO_2000001086a18100",
    "_links": {
      "self": {
        "href": "/api/storage/bridges/2000001086a18100"
      }
    }
  },
  {
    "wwn": "2000001086a18380",
```

```
    "name": "ATTO_2000001086a18380",
    "_links": {
      "self": {
        "href": "/api/storage/bridges/2000001086a18380"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/bridges/"
  }
}
}
----

'''

=== 2) Retrieves a specific bridge from the cluster

==== The following example shows the response of the requested bridge. If
there is no bridge with the requested wwn, an error is returned.

'''

----

# The API:
/api/storage/bridges/{wwn}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/bridges/2000001086a18100" -H
"accept: application/hal+json"

# The response:
{
 "wwn": "2000001086a18100",
 "chassis_throughput_state": "ok",
 "dram_single_bit_error_count": 0,
 "firmware_version": "3.10 007A",
 "ip_address": "10.226.57.178",
 "security_enabled": false,
 "monitoring_enabled": true,
 "model": "FibreBridge 7500N",
 "state": "ok",
```

```
  "managed_by": "in_band",
  "serial_number": "FB7500N102450",
  "symbolic_name": "RTP-FCSAS02-41KK10",
  "vendor": "atto",
  "name": "ATTO_2000001086a18100",
  "last_reboot": {
  "reason": {
    "message": "Reason: \"FirmwareRestart Command\".",
    "code": "39321683"
  },
  "time": "2020-12-09T00:47:58-05:00"
  },
  "paths": [
  {
    "name": "0e",
    "node": {
      "name": "sti8080mcc-htp-005",
      "uuid": "ecc3d992-3a86-11eb-9fab-00a0985a6024",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/ecc3d992-3a86-11eb-9fab-00a0985a6024"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/ports/ecc3d992-3a86-11eb-9fab-
00a0985a6024/0e"
      }
    },
    "target_port": {
      "wwn": "2100001086a18380"
    }
  }
  ],
  "temperature_sensor": {
  "name": "Chassis Temperature Sensor",
  "minimum": 0,
  "maximum": 90,
  "reading": 54,
  "state": "ok"
  },
  "fc_ports": [
{
  "id": 1,
  "state": "online",
```

```
  "enabled": true,
  "negotiated_data_rate": 8,
  "configured_data_rate": 8,
  "conn_mode": "ptp",
  "data_rate_capability": 16,
  "wwn": "2100001086a18100",
  "peer_wwn": "0000000000000000",
  "sfp": {
   "vendor": "FINISAR CORP.",
   "serial_number": "UW106SA",
   "part_number": "FTLF8529P3BCV",
   "data_rate_capability": 16
  }
 },
 {
  "id": 2,
  "state": "online",
  "enabled": true,
  "negotiated_data_rate": 16,
  "configured_data_rate": 16,
  "conn_mode": "ptp",
  "data_rate_capability": 16,
  "wwn": "2200001086a18100",
  "peer_wwn": "0000000000000000",
  "sfp": {
   "vendor": "FINISAR CORP.",
   "serial_number": "UW1072B",
   "part_number": "FTLF8529P3BCV",
   "data_rate_capability": 16
   }
  }
 ],
  "power_supply_units": [
 {
  "name": "A",
  "state": "ok"
 },
 {
  "name": "B",
  "state": "ok"
 }
  ],
  "sas_ports": [
 {
  "id": 1,
   "state": "online",
```

```
    "enabled": true,
    "data_rate_capability": 12,
    "negotiated_data_rate": 6,
    "wwn": "5001086000a18100",
    "phy_1": {
      "state": "online"
    },
    "phy_2": {
      "state": "online"
    },
    "phy_3": {
      "state": "online"
    },
    "phy_4": {
      "state": "online"
    },
    "cable": {
      "vendor": "Molex Inc.",
      "serial_number": "618130935",
      "technology": "Passive Copper 5m ID:00",
      "part_number": "112-00431"
    }
  },
  {
    "state": "offline",
    "enabled": false,
    "data_rate_capability": 12,
    "negotiated_data_rate": 0,
    "wwn": "5001086000a18104",
    "phy_1": {
      "state": "offline"
    },
    "phy_2": {
      "state": "offline"
    },
    "phy_3": {
      "state": "offline"
    },
    "phy_4": {
      "state": "offline"
    },
  },
  {
    "state": "offline",
    "enabled": false,
    "data_rate_capability": 12,
```

```
  "negotiated_data_rate": 0,
  "wwn": "5001086000a18108",
  "phy_1": {
    "state": "offline"
  },
  "phy_2": {
    "state": "offline"
  },
  "phy_3": {
    "state": "offline"
  },
  "phy_4": {
    "state": "offline"
  },
 },
{
  "state": "offline",
  "enabled": false,
  "data_rate_capability": 12,
  "negotiated_data_rate": 0,
  "wwn": "5001086000a1810c",
  "phy_1": {
    "state": "offline"
  },
  "phy_2": {
    "state": "offline"
  },
  "phy_3": {
    "state": "offline"
  },
  "phy_4": {
    "state": "offline"
  },
}
],
"_links": {
"self": {
  "href": "/api/storage/bridges/2000001086a18100"
 }
}
}
----

'''
```

```
[[ID80fc7be525bfaaf79d335c9f54936338]]
= Retrieve a collection of bridges

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/bridges`#

*Introduced In:* 9.9

Retrieves a collection of bridges.

== Related ONTAP commands

* `storage bridge show`

== Learn more

* xref:{relative_path}storage_bridges_endpoint_overview.html[DOC
/storage/bridges]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|sas_ports.state
|string
|query
|False
a|Filter by sas_ports.state


|sas_ports.phy_2.state
|string
|query
|False
a|Filter by sas_ports.phy_2.state
```

```
|sas_ports.phy_1.state
|string
|query
|False
a|Filter by sas_ports.phy_1.state


|sas_ports.id
|integer
|query
|False
a|Filter by sas_ports.id


|sas_ports.wwn
|string
|query
|False
a|Filter by sas_ports.wwn


|sas_ports.enabled
|boolean
|query
|False
a|Filter by sas_ports.enabled


|sas_ports.phy_4.state
|string
|query
|False
a|Filter by sas_ports.phy_4.state


|sas_ports.cable.serial_number
|string
|query
|False
a|Filter by sas_ports.cable.serial_number


|sas_ports.cable.technology
|string
|query
|False
```

```
a|Filter by sas_ports.cable.technology


|sas_ports.cable.part_number
|string
|query
|False
a|Filter by sas_ports.cable.part_number


|sas_ports.cable.vendor
|string
|query
|False
a|Filter by sas_ports.cable.vendor


|sas_ports.negotiated_data_rate
|number
|query
|False
a|Filter by sas_ports.negotiated_data_rate


|sas_ports.data_rate_capability
|number
|query
|False
a|Filter by sas_ports.data_rate_capability


|sas_ports.phy_3.state
|string
|query
|False
a|Filter by sas_ports.phy_3.state


|wwn
|string
|query
|False
a|Filter by wwn


|firmware_version
|string
```

```
|query
|False
a|Filter by firmware_version


|symbolic_name
|string
|query
|False
a|Filter by symbolic_name


|ip_address
|string
|query
|False
a|Filter by ip_address


|serial_number
|string
|query
|False
a|Filter by serial_number


|dram_single_bit_error_count
|integer
|query
|False
a|Filter by dram_single_bit_error_count


|temperature_sensor.minimum
|integer
|query
|False
a|Filter by temperature_sensor.minimum


|temperature_sensor.name
|string
|query
|False
a|Filter by temperature_sensor.name
```

```
|temperature_sensor.state
|string
|query
|False
a|Filter by temperature_sensor.state


|temperature_sensor.maximum
|integer
|query
|False
a|Filter by temperature_sensor.maximum


|temperature_sensor.reading
|integer
|query
|False
a|Filter by temperature_sensor.reading


|fc_ports.peer_wwn
|string
|query
|False
a|Filter by fc_ports.peer_wwn


|fc_ports.state
|string
|query
|False
a|Filter by fc_ports.state


|fc_ports.enabled
|boolean
|query
|False
a|Filter by fc_ports.enabled


|fc_ports.wwn
|string
|query
|False
a|Filter by fc_ports.wwn
```

```
|fc_ports.id
|integer
|query
|False
a|Filter by fc_ports.id


|fc_ports.data_rate_capability
|number
|query
|False
a|Filter by fc_ports.data_rate_capability


|fc_ports.connection_mode
|string
|query
|False
a|Filter by fc_ports.connection_mode


|fc_ports.negotiated_data_rate
|number
|query
|False
a|Filter by fc_ports.negotiated_data_rate


|fc_ports.configured_data_rate
|number
|query
|False
a|Filter by fc_ports.configured_data_rate


|fc_ports.sfp.vendor
|string
|query
|False
a|Filter by fc_ports.sfp.vendor


|fc_ports.sfp.data_rate_capability
|number
|query
```

```
|False
a|Filter by fc_ports.sfp.data_rate_capability


|fc_ports.sfp.serial_number
|string
|query
|False
a|Filter by fc_ports.sfp.serial_number


|fc_ports.sfp.part_number
|string
|query
|False
a|Filter by fc_ports.sfp.part_number


|paths.name
|string
|query
|False
a|Filter by paths.name


|paths.target_port.wwn
|string
|query
|False
a|Filter by paths.target_port.wwn


|paths.target_port.id
|string
|query
|False
a|Filter by paths.target_port.id


|paths.target_port.name
|string
|query
|False
a|Filter by paths.target_port.name


|paths.node.name
```

```
|string
|query
|False
a|Filter by paths.node.name


|paths.node.uuid
|string
|query
|False
a|Filter by paths.node.uuid


|paths.source_port.id
|string
|query
|False
a|Filter by paths.source_port.id


|paths.source_port.name
|string
|query
|False
a|Filter by paths.source_port.name


|vendor
|string
|query
|False
a|Filter by vendor


|power_supply_units.state
|string
|query
|False
a|Filter by power_supply_units.state


|power_supply_units.name
|string
|query
|False
a|Filter by power_supply_units.name
```

```
|security_enabled
|boolean
|query
|False
a|Filter by security_enabled


|name
|string
|query
|False
a|Filter by name


|state
|string
|query
|False
a|Filter by state


|errors.type
|string
|query
|False
a|Filter by errors.type


|errors.reason.arguments.message
|string
|query
|False
a|Filter by errors.reason.arguments.message

* Introduced in: 9.10


|errors.reason.arguments.code
|string
|query
|False
a|Filter by errors.reason.arguments.code

* Introduced in: 9.10
```

```
|errors.reason.code
|string
|query
|False
a|Filter by errors.reason.code


|errors.reason.message
|string
|query
|False
a|Filter by errors.reason.message


|errors.severity
|string
|query
|False
a|Filter by errors.severity


|errors.component.id
|integer
|query
|False
a|Filter by errors.component.id


|errors.component.name
|string
|query
|False
a|Filter by errors.component.name


|errors.component.unique_id
|string
|query
|False
a|Filter by errors.component.unique_id


|monitoring_enabled
|boolean
|query
|False
a|Filter by monitoring_enabled
```

```
|model
|string
|query
|False
a|Filter by model


|last_reboot.reason.arguments.message
|string
|query
|False
a|Filter by last_reboot.reason.arguments.message

* Introduced in: 9.10


|last_reboot.reason.arguments.code
|string
|query
|False
a|Filter by last_reboot.reason.arguments.code

* Introduced in: 9.10


|last_reboot.reason.code
|string
|query
|False
a|Filter by last_reboot.reason.code


|last_reboot.reason.message
|string
|query
|False
a|Filter by last_reboot.reason.message


|last_reboot.time
|string
|query
|False
a|Filter by last_reboot.time
```

```
|managed_by
|string
|query
|False
a|Filter by managed_by


|chassis_throughput_state
|string
|query
|False
a|Filter by chassis_throughput_state


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#storage_bridge[storage_bridge]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "chassis_throughput_state": "ok",
      "errors": [
        {
          "component": {
            "name": "string",
            "unique_id": "string"
          },
          "reason": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          },
          "severity": "string",
          "type": "string"
        }
      ],
      "fc_ports": [
        {
          "configured_data_rate": 0,
          "connection_mode": "string",
          "data_rate_capability": 2,
          "negotiated_data_rate": 0,
          "peer_wwn": "200650eb1a238892",
          "sfp": {
            "data_rate_capability": 2,
            "part_number": "string",
            "serial_number": "string",
            "vendor": "string"
          },
          "state": "string",
```

```
        "wwn": "2100001086a54100"
      }
    ],
    "firmware_version": "4.10 007A",
    "ip_address": "string",
    "last_reboot": {
      "reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "time": "2020-12-09 00:47:58 -0500"
    },
    "managed_by": "string",
    "model": "FibreBridge6500N",
    "name": "ATTO_FibreBridge6500N_1",
    "paths": [
      {
        "name": "2c",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "source_port": {
          "id": "100050eb1a238892",
          "name": "rtp-fc03-41kk11:1"
        },
        "target_port": {
          "id": "100050eb1a238892",
          "name": "rtp-fc03-41kk11:6",
          "wwn": "2100001086a54100"
        }
      }
    ],
    "power_supply_units": [
      {
```

```
          "name": "string",
          "state": "string"
        }
      ],
      "sas_ports": [
        {
          "cable": {
            "part_number": "string",
            "serial_number": "string",
            "technology": "string",
            "vendor": "string"
          },
          "data_rate_capability": 0,
          "negotiated_data_rate": 0,
          "phy_1": {
            "state": "string"
          },
          "phy_2": {
            "state": "string"
          },
          "phy_3": {
            "state": "string"
          },
          "phy_4": {
            "state": "string"
          },
          "state": "string",
          "wwn": "2100001086a54100"
        }
      ],
      "serial_number": "FB7600N100004",
      "state": "string",
      "symbolic_name": "rtp-fcsas03-41kk11",
      "temperature_sensor": {
        "name": "Chassis temperature sensor",
        "state": "string"
      },
      "vendor": "atto",
      "wwn": 2000001086600476
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#component]
[.api-collapsible-fifth-title]
component

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Bridge error component ID
```

```
|name
|string
a|Bridge error component name


|unique_id
|string
a|Bridge error component unique ID


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|component
|link:#component[component]
a|

|reason
|link:#error[error]
a|

|severity
|string
a|Bridge error severity


|type
|string
a|Bridge error type


|===
```

```
[#sfp]
[.api-collapsible-fifth-title]
sfp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|data_rate_capability
|number
a|Bridge FC port SFP data rate capability, in Gbps


|part_number
|string
a|

|serial_number
|string
a|Bridge FC port SFP serial number


|vendor
|string
a|Bridge FC port SFP vendor


|===


[#fc_ports]
[.api-collapsible-fifth-title]
fc_ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|configured_data_rate
|number
a|Bridge FC port configured data rate, in Gbps
```

```
|connection_mode
|string
a|Bridge FC port configured connection mode


|data_rate_capability
|number
a|Bridge FC port data rate capability, in Gbps


|enabled
|boolean
a|Indicates whether the bridge FC port is enabled.


|id
|integer
a|Bridge FC port index


|negotiated_data_rate
|number
a|Bridge FC port negotiated data rate, in Gbps


|peer_wwn
|string
a|Bridge FC port peer port world wide name


|sfp
|link:#sfp[sfp]
a|

|state
|string
a|Bridge FC port state


|wwn
|string
a|Bridge FC port world wide name


|===
```

```
[#error]
[.api-collapsible-fifth-title]
error

The error message and code explaining why the bridge rebooted.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#last_reboot]
[.api-collapsible-fifth-title]
last_reboot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|The error message and code explaining why the bridge rebooted.
```

```
|time
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#source_port]
[.api-collapsible-fifth-title]
source_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Initiator side switch port id


|name
|string
a|Initiator side switch port name


|===


[#target_port]
[.api-collapsible-fifth-title]
target_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Target side switch port id


|name
|string
a|Target side switch port name


|wwn
```

```
|string
a|Target side switch port world wide name


|===


[#paths]
[.api-collapsible-fifth-title]
paths

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|source_port
|link:#source_port[source_port]
a|

|target_port
|link:#target_port[target_port]
a|

|===


[#power_supply_units]
[.api-collapsible-fifth-title]
power_supply_units

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|Power supply unit name


|state
|string
a|Power supply unit state


|===


[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|part_number
|string
a|Bridge cable part number


|serial_number
|string
a|Bridge cable serial number


|technology
|string
a|Bridge cable type


|vendor
|string
a|Bridge cable vendor


|===


[#phy_1]
[.api-collapsible-fifth-title]
```

```
phy_1

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY1 state


|===


[#phy_2]
[.api-collapsible-fifth-title]
phy_2

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY2 state


|===


[#phy_3]
[.api-collapsible-fifth-title]
phy_3

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY3 state
```

```
|===


[#phy_4]
[.api-collapsible-fifth-title]
phy_4

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY4 state


|===


[#sas_ports]
[.api-collapsible-fifth-title]
sas_ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cable
|link:#cable[cable]
a|

|data_rate_capability
|number
a|Bridge SAS port data rate capability, in Gbps


|enabled
|boolean
a|Indicates whether a bridge SAS port is enabled.
```

```
|id
|integer
a|Bridge SAS port index


|negotiated_data_rate
|number
a|Bridge SAS port negotiated data rate, in Gbps


|phy_1
|link:#phy_1[phy_1]
a|

|phy_2
|link:#phy_2[phy_2]
a|

|phy_3
|link:#phy_3[phy_3]
a|

|phy_4
|link:#phy_4[phy_4]
a|

|state
|string
a|Bridge SAS port state


|wwn
|string
a|Bridge SAS port world wide name


|===


[#temperature_sensor]
[.api-collapsible-fifth-title]
temperature_sensor

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|maximum
|integer
a|Maximum safe operating temperature, in degrees Celsius.


|minimum
|integer
a|Minimum safe operating temperature, in degrees Celsius.


|name
|string
a|Temperature sensor name


|reading
|integer
a|Chassis temperature sensor reading, in degrees Celsius.


|state
|string
a|

|===


[#storage_bridge]
[.api-collapsible-fifth-title]
storage_bridge

[cols=3*,options=header]
|===
|Name
|Type
|Description

|chassis_throughput_state
|string
a|Chassis throughput status


|dram_single_bit_error_count
|integer
a|
```

```
|errors
|array[link:#errors[errors]]
a|

|fc_ports
|array[link:#fc_ports[fc_ports]]
a|

|firmware_version
|string
a|Bridge firmware version


|ip_address
|string
a|IP Address


|last_reboot
|link:#last_reboot[last_reboot]
a|

|managed_by
|string
a|

|model
|string
a|Bridge model


|monitoring_enabled
|boolean
a|Indicates whether monitoring is enabled for the bridge.


|name
|string
a|Bridge name


|paths
|array[link:#paths[paths]]
a|

|power_supply_units
```

```
|array[link:#power_supply_units[power_supply_units]]
a|

|sas_ports
|array[link:#sas_ports[sas_ports]]
a|

|security_enabled
|boolean
a|Indicates whether security is enabled for the bridge.


|serial_number
|string
a|Bridge serial number


|state
|string
a|Bridge state


|symbolic_name
|string
a|Bridge symbolic name


|temperature_sensor
|link:#temperature_sensor[temperature_sensor]
a|

|vendor
|string
a|Bridge vendor


|wwn
|string
a|Bridge world wide name


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4becef5370dcba2a7c2ae307e661ee97]]
= Retrieve a specific bridge

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/bridges/{wwn}`#

*Introduced In:* 9.9

Retrieves a specific bridge

== Related ONTAP commands
```

```
* `storage bridge show`

== Learn more

* xref:{relative_path}storage_bridges_endpoint_overview.html[DOC
/storage/bridges]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|wwn
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|chassis_throughput_state
|string
a|Chassis throughput status
```

```
|dram_single_bit_error_count
|integer
a|

|errors
|array[link:#errors[errors]]
a|

|fc_ports
|array[link:#fc_ports[fc_ports]]
a|

|firmware_version
|string
a|Bridge firmware version


|ip_address
|string
a|IP Address


|last_reboot
|link:#last_reboot[last_reboot]
a|

|managed_by
|string
a|

|model
|string
a|Bridge model


|monitoring_enabled
|boolean
a|Indicates whether monitoring is enabled for the bridge.


|name
|string
a|Bridge name
```

```
|paths
|array[link:#paths[paths]]
a|

|power_supply_units
|array[link:#power_supply_units[power_supply_units]]
a|

|sas_ports
|array[link:#sas_ports[sas_ports]]
a|

|security_enabled
|boolean
a|Indicates whether security is enabled for the bridge.


|serial_number
|string
a|Bridge serial number


|state
|string
a|Bridge state


|symbolic_name
|string
a|Bridge symbolic name


|temperature_sensor
|link:#temperature_sensor[temperature_sensor]
a|

|vendor
|string
a|Bridge vendor


|wwn
|string
a|Bridge world wide name


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "chassis_throughput_state": "ok",
  "errors": [
    {
      "component": {
        "name": "string",
        "unique_id": "string"
      },
      "reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "severity": "string",
      "type": "string"
    }
  ],
  "fc_ports": [
    {
      "configured_data_rate": 0,
      "connection_mode": "string",
      "data_rate_capability": 2,
      "negotiated_data_rate": 0,
      "peer_wwn": "200650eb1a238892",
      "sfp": {
        "data_rate_capability": 2,
        "part_number": "string",
        "serial_number": "string",
        "vendor": "string"
      },
      "state": "string",
      "wwn": "2100001086a54100"
    }
  ],
  "firmware_version": "4.10 007A",
```

```
  "ip_address": "string",
  "last_reboot": {
    "reason": {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist"
    },
    "time": "2020-12-09 00:47:58 -0500"
  },
  "managed_by": "string",
  "model": "FibreBridge6500N",
  "name": "ATTO_FibreBridge6500N_1",
  "paths": [
    {
      "name": "2c",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "source_port": {
        "id": "100050eb1a238892",
        "name": "rtp-fc03-41kk11:1"
      },
      "target_port": {
        "id": "100050eb1a238892",
        "name": "rtp-fc03-41kk11:6",
        "wwn": "2100001086a54100"
      }
    }
  ],
  "power_supply_units": [
    {
      "name": "string",
      "state": "string"
    }
  ],
```

```
  "sas_ports": [
    {
      "cable": {
        "part_number": "string",
        "serial_number": "string",
        "technology": "string",
        "vendor": "string"
      },
      "data_rate_capability": 0,
      "negotiated_data_rate": 0,
      "phy_1": {
        "state": "string"
      },
      "phy_2": {
        "state": "string"
      },
      "phy_3": {
        "state": "string"
      },
      "phy_4": {
        "state": "string"
      },
      "state": "string",
      "wwn": "2100001086a54100"
    }
  ],
  "serial_number": "FB7600N100004",
  "state": "string",
  "symbolic_name": "rtp-fcsas03-41kk11",
  "temperature_sensor": {
    "name": "Chassis temperature sensor",
    "state": "string"
  },
  "vendor": "atto",
  "wwn": 2000001086600476
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#component]
[.api-collapsible-fifth-title]
component

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|id
|integer
a|Bridge error component ID


|name
|string
a|Bridge error component name


|unique_id
|string
a|Bridge error component unique ID


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|component
|link:#component[component]
a|

|reason
|link:#error[error]
a|

|severity
|string
a|Bridge error severity


|type
|string
```

```
a|Bridge error type


|===


[#sfp]
[.api-collapsible-fifth-title]
sfp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|data_rate_capability
|number
a|Bridge FC port SFP data rate capability, in Gbps


|part_number
|string
a|

|serial_number
|string
a|Bridge FC port SFP serial number


|vendor
|string
a|Bridge FC port SFP vendor


|===


[#fc_ports]
[.api-collapsible-fifth-title]
fc_ports

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|configured_data_rate
|number
a|Bridge FC port configured data rate, in Gbps


|connection_mode
|string
a|Bridge FC port configured connection mode


|data_rate_capability
|number
a|Bridge FC port data rate capability, in Gbps


|enabled
|boolean
a|Indicates whether the bridge FC port is enabled.


|id
|integer
a|Bridge FC port index


|negotiated_data_rate
|number
a|Bridge FC port negotiated data rate, in Gbps


|peer_wwn
|string
a|Bridge FC port peer port world wide name


|sfp
|link:#sfp[sfp]
a|

|state
|string
a|Bridge FC port state


|wwn
|string
```

```
a|Bridge FC port world wide name


|===


[#error]
[.api-collapsible-fifth-title]
error

The error message and code explaining why the bridge rebooted.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#last_reboot]
[.api-collapsible-fifth-title]
last_reboot

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|reason
|link:#error[error]
a|The error message and code explaining why the bridge rebooted.


|time
|string
a|

|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
```

```
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#source_port]
[.api-collapsible-fifth-title]
source_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Initiator side switch port id


|name
|string
a|Initiator side switch port name


|===
```

```
[#target_port]
[.api-collapsible-fifth-title]
target_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Target side switch port id


|name
|string
a|Target side switch port name


|wwn
|string
a|Target side switch port world wide name


|===


[#paths]
[.api-collapsible-fifth-title]
paths

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|source_port
```

```
|link:#source_port[source_port]
a|

|target_port
|link:#target_port[target_port]
a|

|===


[#power_supply_units]
[.api-collapsible-fifth-title]
power_supply_units

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Power supply unit name


|state
|string
a|Power supply unit state


|===


[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|part_number
|string
a|Bridge cable part number
```

```
|serial_number
|string
a|Bridge cable serial number


|technology
|string
a|Bridge cable type


|vendor
|string
a|Bridge cable vendor


|===


[#phy_1]
[.api-collapsible-fifth-title]
phy_1

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY1 state


|===


[#phy_2]
[.api-collapsible-fifth-title]
phy_2

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|state
|string
a|Bridge SAS port PHY2 state


|===


[#phy_3]
[.api-collapsible-fifth-title]
phy_3

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY3 state


|===


[#phy_4]
[.api-collapsible-fifth-title]
phy_4

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|Bridge SAS port PHY4 state


|===


[#sas_ports]
[.api-collapsible-fifth-title]
sas_ports
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|cable
|link:#cable[cable]
a|

|data_rate_capability
|number
a|Bridge SAS port data rate capability, in Gbps


|enabled
|boolean
a|Indicates whether a bridge SAS port is enabled.


|id
|integer
a|Bridge SAS port index


|negotiated_data_rate
|number
a|Bridge SAS port negotiated data rate, in Gbps


|phy_1
|link:#phy_1[phy_1]
a|

|phy_2
|link:#phy_2[phy_2]
a|

|phy_3
|link:#phy_3[phy_3]
a|

|phy_4
|link:#phy_4[phy_4]
a|
```

```
|state
|string
a|Bridge SAS port state


|wwn
|string
a|Bridge SAS port world wide name


|===


[#temperature_sensor]
[.api-collapsible-fifth-title]
temperature_sensor

[cols=3*,options=header]
|===
|Name
|Type
|Description

|maximum
|integer
a|Maximum safe operating temperature, in degrees Celsius.


|minimum
|integer
a|Minimum safe operating temperature, in degrees Celsius.


|name
|string
a|Temperature sensor name


|reading
|integer
a|Chassis temperature sensor reading, in degrees Celsius.


|state
|string
a|
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage cluster-wide storage

:leveloffset: +1
```

```
[[ID7af5a4dfceeca4cb97be938937231273]]
= Manage cluster-wide storage



Retrieves cluster-wide storage details across the different tiers.
Storage details include storage efficiency, block storage and cloud
storage information.

'''

Example

=== Retrieving cluster-wide storage details

The following example shows the details returned for a GET request on
cluster-wide storage:

'''

----

# The API:
/api/storage/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/cluster" -H "accept:
application/hal+json"

# The response:
{
"efficiency": {
  "savings": 143360,
  "ratio": 1.134099616858238,
  "logical_used": 1212416
},
"efficiency_without_snapshots": {
  "savings": 0,
  "ratio": 1,
  "logical_used": 167936
},
"efficiency_without_snapshots_flexclones": {
  "savings": 0,
  "ratio": 1,
  "logical_used": 167936
```

```
},
"block_storage": {
  "used": 6269812736,
  "size": 56125612032,
  "available": 49855799296,
  "physical_used": 1838284800,
  "inactive_data": 0,
  "medias": [
    {
      "type": "ssd",
      "size": 9891430400,
      "available": 3728039936,
      "used": 6163390464,
      "physical_used": 1832886272,
      "efficiency": {
        "savings": 0,
        "ratio": 1,
        "logical_used": 0
      },
      "efficiency_without_snapshots": {
        "savings": 0,
        "ratio": 1,
        "logical_used": 0
      },
      "efficiency_without_snapshots_flexclones": {
        "savings": 0,
        "ratio": 1,
        "logical_used": 0
      }
    },
    {
      "type": "vmdisk",
      "size": 46234181632,
      "available": 46127759360,
      "used": 106422272,
      "physical_used": 5398528,
      "efficiency": {
        "savings": 282624,
        "ratio": 1.303964757709251,
        "logical_used": 1212416
      },
      "efficiency_without_snapshots": {
        "savings": 0,
        "ratio": 1,
        "logical_used": 167936
      },
```

```
      "efficiency_without_snapshots_flexclones": {
        "savings": 0,
        "ratio": 1,
        "logical_used": 167936
      }
    }
  ]
}
}
----

'''




[[IDe4f3a00364269b2b8abec77ae1397d85]]
= Report cluster-wide storage details across different tiers

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/cluster`#

*Introduced In:* 9.6

Retrieves cluster-wide storage details across the different tiers. By
default, this endpoint returns all fields.
Storage details include storage efficiency, block storage and cloud
storage information.
Supports the following roles: admin, and readonly.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|efficiency_without_snapshots.savings
|integer
|query
|False
```

```
a|Filter by efficiency_without_snapshots.savings

* Introduced in: 9.14


|efficiency_without_snapshots.logical_used
|integer
|query
|False
a|Filter by efficiency_without_snapshots.logical_used

* Introduced in: 9.14


|efficiency_without_snapshots.ratio
|number
|query
|False
a|Filter by efficiency_without_snapshots.ratio

* Introduced in: 9.14


|efficiency.savings
|integer
|query
|False
a|Filter by efficiency.savings

* Introduced in: 9.14


|efficiency.logical_used
|integer
|query
|False
a|Filter by efficiency.logical_used

* Introduced in: 9.14


|efficiency.ratio
|number
|query
|False
a|Filter by efficiency.ratio
```

```
* Introduced in: 9.14


|block_storage.medias.size
|integer
|query
|False
a|Filter by block_storage.medias.size

* Introduced in: 9.14


|block_storage.medias.available
|integer
|query
|False
a|Filter by block_storage.medias.available

* Introduced in: 9.14


|block_storage.medias.efficiency_without_snapshots_flexclones.savings
|integer
|query
|False
a|Filter by
block_storage.medias.efficiency_without_snapshots_flexclones.savings

* Introduced in: 9.14


|block_storage.medias.efficiency_without_snapshots_flexclones.logical_used
|integer
|query
|False
a|Filter by
block_storage.medias.efficiency_without_snapshots_flexclones.logical_used

* Introduced in: 9.14


|block_storage.medias.efficiency_without_snapshots_flexclones.ratio
|number
|query
|False
a|Filter by
block_storage.medias.efficiency_without_snapshots_flexclones.ratio
```

```
* Introduced in: 9.14


|block_storage.medias.efficiency.savings
|integer
|query
|False
a|Filter by block_storage.medias.efficiency.savings

* Introduced in: 9.14


|block_storage.medias.efficiency.logical_used
|integer
|query
|False
a|Filter by block_storage.medias.efficiency.logical_used

* Introduced in: 9.14


|block_storage.medias.efficiency.ratio
|number
|query
|False
a|Filter by block_storage.medias.efficiency.ratio

* Introduced in: 9.14


|block_storage.medias.efficiency_without_snapshots.savings
|integer
|query
|False
a|Filter by block_storage.medias.efficiency_without_snapshots.savings

* Introduced in: 9.14


|block_storage.medias.efficiency_without_snapshots.logical_used
|integer
|query
|False
a|Filter by block_storage.medias.efficiency_without_snapshots.logical_used

* Introduced in: 9.14
```

```
|block_storage.medias.efficiency_without_snapshots.ratio
|number
|query
|False
a|Filter by block_storage.medias.efficiency_without_snapshots.ratio

* Introduced in: 9.14


|block_storage.medias.physical_used
|integer
|query
|False
a|Filter by block_storage.medias.physical_used

* Introduced in: 9.14


|block_storage.medias.type
|string
|query
|False
a|Filter by block_storage.medias.type

* Introduced in: 9.14


|block_storage.medias.used
|integer
|query
|False
a|Filter by block_storage.medias.used

* Introduced in: 9.14


|block_storage.size
|integer
|query
|False
a|Filter by block_storage.size

* Introduced in: 9.14
```

```
|block_storage.inactive_data
|integer
|query
|False
a|Filter by block_storage.inactive_data

* Introduced in: 9.14


|block_storage.available
|integer
|query
|False
a|Filter by block_storage.available

* Introduced in: 9.14


|block_storage.physical_used
|integer
|query
|False
a|Filter by block_storage.physical_used

* Introduced in: 9.14


|block_storage.used
|integer
|query
|False
a|Filter by block_storage.used

* Introduced in: 9.14


|cloud_storage.used
|integer
|query
|False
a|Filter by cloud_storage.used

* Introduced in: 9.14


|efficiency_without_snapshots_flexclones.savings
|integer
```

```
|query
|False
a|Filter by efficiency_without_snapshots_flexclones.savings

* Introduced in: 9.14


|efficiency_without_snapshots_flexclones.logical_used
|integer
|query
|False
a|Filter by efficiency_without_snapshots_flexclones.logical_used

* Introduced in: 9.14


|efficiency_without_snapshots_flexclones.ratio
|number
|query
|False
a|Filter by efficiency_without_snapshots_flexclones.ratio

* Introduced in: 9.14


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_storage
|link:#block_storage[block_storage]
a|Configuration information for the locally attached portion of the
```

```
storage across the cluster. When a cloud store is also used by the
storage, this is referred to as the performance tier.


|cloud_storage
|link:#cloud_storage[cloud_storage]
a|Configuration information for the cloud storage portion of all the
aggregates across the cluster. This is referred to as the capacity tier.


|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and FlexClone volumes.


|metric
|link:#metric[metric]
a|Cluster capacity numbers, such as total size, used size, and available
size.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "block_storage": {
    "available": 0,
    "inactive_data": 0,
    "medias": [
      {
```

```
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots_flexclones": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "type": "string"
      }
    ],
    "physical_used": 0,
    "size": 0,
    "used": 0
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#efficiency]
[.api-collapsible-fifth-title]
```

```
efficiency

Storage efficiency.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots]
[.api-collapsible-fifth-title]
efficiency_without_snapshots

Storage efficiency that does not include the savings provided by
snapshots.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used
```

```
|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#efficiency_without_snapshots_flexclones]
[.api-collapsible-fifth-title]
efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by snapshots
and FlexClone volumes.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|logical_used
|integer
a|Logical used


|ratio
|number
a|Data reduction ratio (logical_used / used)


|savings
|integer
a|Space saved by storage efficiencies (logical_used - used)


|===


[#medias]
```

```
[.api-collapsible-fifth-title]
medias

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|Available space across the cluster based on media type.


|efficiency
|link:#efficiency[efficiency]
a|Storage efficiency.


|efficiency_without_snapshots
|link:#efficiency_without_snapshots[efficiency_without_snapshots]
a|Storage efficiency that does not include the savings provided by
snapshots.


|efficiency_without_snapshots_flexclones
|link:#efficiency_without_snapshots_flexclones[efficiency_without_snapshot
s_flexclones]
a|Storage efficiency that does not include the savings provided by
snapshots and FlexClone volumes.


|physical_used
|integer
a|Total physical used space across the cluster based on media type.


|size
|integer
a|Total space across the cluster based on media type.


|type
|string
a|The type of media being used.
```

```
|used
|integer
a|Used space across the cluster based on media type.


|===


[#block_storage]
[.api-collapsible-fifth-title]
block_storage

Configuration information for the locally attached portion of the storage
across the cluster. When a cloud store is also used by the storage, this
is referred to as the performance tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|Available space across the cluster.


|inactive_data
|integer
a|Inactive data across the cluster.


|medias
|array[link:#medias[medias]]
a|Configuration information based on type of media. For example, SSD media
type information includes the sum of all the SSD storage across the
cluster.


|physical_used
|integer
a|Total physical used space across the cluster.


|size
|integer
```

```
a|Total space across the cluster.


|used
|integer
a|Used space (includes volume reserves) across the cluster.


|===


[#cloud_storage]
[.api-collapsible-fifth-title]
cloud_storage

Configuration information for the cloud storage portion of all the
aggregates across the cluster. This is referred to as the capacity tier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|used
|integer
a|Total space used in cloud.


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#metric]
[.api-collapsible-fifth-title]
metric

Cluster capacity numbers, such as total size, used size, and available
size.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage disks

:leveloffset: +1


[[IDaee7140e483d460ab4e8537a22b78263]]
= Manage disks



== Retrieving storage disk information

The storage disk GET API retrieves all of the disks in the cluster.

'''

== Examples

=== 1) Retrieve a list of disks from the cluster.

==== The following example shows the response with a list of disks in the
cluster:

'''

----

# The API:
/api/storage/disks

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks" -H "accept:
application/hal+json"

# The response:
{
```

```
"records": [
  {
    "name": "1.24.4",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.4"
      }
    }
  },
  {
    "name": "1.24.3",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.3"
      }
    }
  },
  {
    "name": "1.24.5",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.5"
      }
    }
  },
  {
    "name": "1.24.0",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.0"
      }
    }
  },
  {
    "name": "1.24.2",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.2"
      }
    }
  },
  {
    "name": "1.24.1",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.1"
```

```
      }
    }
  }
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/storage/disks"
  }
}
}
----

'''

=== 2) Retrieve a specific disk from the cluster.

==== The following example shows the response of the requested disk. If
there is no disk with the requested name, an error is returned:

'''

----

# The API:
/api/storage/disks/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks/1.24.3" -H "accept:
application/hal+json"

# The response:
{
"name": "1.24.3",
"uid":
"50000394:0808AA88:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000",
"serial_number": "EC47PC5021SW",
"model": "X421_FAL12450A10",
"vendor": "NETAPP",
"firmware_version": "NA02",
"usable_size": 438304768000,
"rpm": 10000,
"type": "sas",
"effective_type": "sas",
"class": "performance",
```

```
"container_type": "aggregate",
"pool": "pool0",
"state": "present",
"node": {
  "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
  "name": "node-2",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
    }
  }
},
"home_node": {
  "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
  "name": "node-2",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
    }
  }
},
"aggregates": [
  {
    "uuid": "3fd9c345-ba91-4949-a7b1-6e2b898d74e3",
    "name": "node_2_SAS_1",
    "_links": {
      "self": {
        "href": "/api/storage/aggregates/3fd9c345-ba91-4949-a7b1-
6e2b898d74e3"
      }
    }
  }
],
"shelf": {
  "uid": "10318311901725526608",
  "_links": {
    "self": {
      "href": "/api/storage/shelves/10318311901725526608"
    }
  }
},
"local": true,
"paths": [
{
   "initiator": "3a",
   "port_name": "B",
```

```
    "port_type": "sas",
    "wwnn": "5000cca02f0e6768",
    "wwpn": "5000cca02f0e676a",
    "node": {
      "name": "vsim3",
      "uuid": "3f7fa09a-5c56-11ec-b366-005056bbbb3f",
      "_links": {
         "self": {
           "href": "/api/cluster/nodes/3f7fa09a-5c56-11ec-b366-
005056bbbb3f"
         }
      }
    }
  },
  {
    "initiator": "3d",
    "port_name": "A",
    "port_type": "sas",
    "wwnn": "5000cca02f0e6768",
    "wwpn": "5000cca02f0e6769",
    "node": {
      "name": "vsim4",
      "uuid": "4f7fa09a-5c56-11ec-b366-005056bbbb3f",
      "_links": {
         "self": {
           "href": "/api/cluster/nodes/4f7fa09a-5c56-11ec-b366-
005056bbbb3f"
         }
      }
    }
},
{
    "initiator": "3d",
    "port_name": "A",
    "port_type": "sas",
    "wwnn": "5000cca02f0e6768",
    "wwpn": "5000cca02f0e6769",
    "node": {
      "name": "vsim3",
      "uuid": "3f7fa09a-5c56-11ec-b366-005056bbbb3f",
      "_links": {
         "self": {
           "href": "/api/cluster/nodes/3f7fa09a-5c56-11ec-b366-
005056bbbb3f"
         }
      }
```

```
    }
  },
  {
    "initiator": "3a",
    "port_name": "B",
    "port_type": "sas",
    "wwnn": "5000cca02f0e6768",
    "wwpn": "5000cca02f0e676a",
    "node": {
      "name": "vsim4",
      "uuid": "4f7fa09a-5c56-11ec-b366-005056bbbb3f",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4f7fa09a-5c56-11ec-b366-
005056bbbb3f"
        }
      }
    }
  }
  ],
  "outage": {
    "persistently_failed": true,
    "reason": {
      "message": "Failed disk. Reason: \"admin failed\".",
      "code": "721081"
    }
  },
  "bay": 3,
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.3"
    }
  },
  "error": [
  {
    "reason": {
      "message": "\"The node is configured with All-Flash Optimized
personality and this disk is not an SSD. The disk needs to be removed from
the system.\"",
      "code": "721082"
    },
    "type": "notallflashdisk"
  }
  ],
  "bytes_per_sector": 512,
  "sector_count": 1172123568,
```

```
"right_size_sector_count": 5579776,
"physical_size": 438804988000,
"stats": {
  "average_latency": 6,
  "throughput": 1957888,
  "iops_total": 12854,
  "path_error_count": 0,
  "power_on_hours": 11797
 }
}
----

'''

=== 3) Retrieving a specific disk from the hypervisor

==== The following example shows the response of the requested disk. If
there is no disk with the requested name, an error is returned:

'''

----

# The API:
/api/storage/disks/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks/NET-3.2" -H "accept:
application/hal+json"

# The response:
{
"name": "NET-3.2",
"uid":
"32343637:65386464:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000",
"serial_number": "3234363765386464",
"model": "PHA-DISK",
"vendor": "NETAPP",
"firmware_version": "0001",
"type": "vmdisk",
"class": "virtual",
"container_type": "mediator",
"pool": "pool0",
"node": {
    "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
```

```
    "name": "example_node_name",
    "_links": {
      "self": {
          "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
      }
    }
},
"home_node": {
    "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
    "name": "example_node_name",
    "_links": {
      "self": {
          "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
      }
    }
},
"local": true,
"paths": [
    {
      "initiator": "0f",
      "port_name": "A",
      "port_type": "sas",
      "wwnn": "53059d50444f5476",
      "wwpn": "53059d50444f5476",
      "vmdisk_hypervisor_file_name": "LUN 4.0",
      "node": {
          "name": "example_node_name",
          "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
          "_links": {
            "self": {
                "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
            }
          }
      }
    },
    {
      "initiator": "0f",
      "port_name": "A",
      "port_type": "sas",
      "wwnn": "53059d50444f5476",
      "wwpn": "53059d50444f5476",
      "vmdisk_hypervisor_file_name": "LUN 2.0",
      "node": {
```

```
          "name": "example_node_name",
          "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
          "_links": {
            "self": {
                "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
            }
          }
      }
    }
],
"outage": {
    "persistently_failed": false,
    "reason": {
      "message": "Failed disk. Reason: \"\".",
      "code": "721081"
    }
},
"self_encrypting": false,
"fips_certified": false,
"bytes_per_sector": 512,
"sector_count": 204808,
"right_size_sector_count": 5579776,
"physical_size": 204808,
"stats": {
    "average_latency": 2157188883,
    "throughput": 4096,
    "iops_total": 1,
    "path_error_count": 0,
    "power_on_hours": 0
},
"_links": {
    "self": {
      "href": "/api/storage/disks/NET-3.2"
    }
}
}
----

'''

== Modifying storage disk

The storage disk PATCH API modifies disk ownership, unfails a disk,
updates encrypting drive authentication keys (AKs), sanitizes encrypting
drives, or sanitizes non-encrypting spare drives in the cluster.
```

```
The storage disk API currently supports patching one attribute at a time.

=== Updating the disk ownership for a specified disk. Disk ownership
cannot be updated on the ASA r2 platform.

=== 1. When the disk is not assigned

When the disk is a spare (or unowned) disk and node name is specified, the
PATCH operation assigns the disk to the specified node.
Optionally, pool name can also be specified along with node name. Accepted
pool names are: pool0, pool1.

=== 2. When the disk is already assigned

When the disk is already assigned (already has a owner), and a new node is
specified, the PATCH operation changes the ownership to the new node.
Optionally, pool name can also be specified along with node name. Accepted
pool names are: pool0, pool1.

=== Removing the disk ownership for a specified disk

When the disk is already assigned, and node name is specified as null (no-
quotes), the PATCH operation removes the owner.

'''

== Examples

=== 1. Update the disk ownership for an unowned disk

'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}}'

# The response:
{
}
----
```

```
'''

=== 2. Update the disk ownership for an already owned disk

'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}}'

# The response:
{
}
----

'''

=== 3. Update the disk pool for a disk (can be either owned or unowned).

'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}, "pool": "pool0"}'

# The response:
{
}
----

'''

=== 4. Rekey the data authentication key (AK) of all encrypting drives to
an authentication key (AK) selected automatically by the system
```

```
'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/storage/disks?name=*&encryption_operation=rekey_data_auto_id" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json"

# The response contains the number of disks attempted.
{
 "num_records": 32
}
----

'''

=== 5. Cryptographically sanitize a spare or broken disk

'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-
name>&encryption_operation=sanitize_disk" -H "accept:
application/hal+json" -H "Content-Type: application/hal+json"

# The response contains the number of disks attempted.
{
 "num_records": 1
}
----

'''

=== 6. Unfailing a disk to a spare.

'''

----
```

```
# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -d
'{"state": "spare"}' -H "accept: application/hal+json" -H "Content-Type:
application/hal+json"

# The response:
{
}
----

'''

=== 7. Unfailing a disk and attempting to reassimilate filesystem labels.

=== If unable or unnecessary to reassimilate filesystem labels, the disk
will be set as spare.

'''

----

# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -d
'{"state": "present"}' -H "accept: application/hal+json" -H "Content-Type:
application/hal+json"

# The response:
{
}
----

'''

=== 8. Sanitize spare disks (non-cryptographically)

'''

----

# The API:
```

```
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -d
'{"sanitize_spare": "true"}' -H "accept: application/hal+json" -H
"Content-Type: application/hal+json"

# The response:
{
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/disks?name=<disk-name>"
  }
}
}
----

'''





[[ID5603481620e799a6752c7ae6547b362f]]
= Retrieve a collection of disks

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/disks`#

*Introduced In:* 9.6

Retrieves a collection of disks.

== Related ONTAP commands

* `storage disk show`

== Learn more

* xref:{relative_path}storage_disks_endpoint_overview.html[DOC
/storage/disks]


== Parameters

[cols=5*,options=header]
```

```
|===

|Name
|Type
|In
|Required
|Description

|bay
|integer
|query
|False
a|Filter by bay


|sector_count
|integer
|query
|False
a|Filter by sector_count

* Introduced in: 9.9


|rated_life_used_percent
|integer
|query
|False
a|Filter by rated_life_used_percent


|protection_mode
|string
|query
|False
a|Filter by protection_mode

* Introduced in: 9.7


|control_standard
|string
|query
|False
a|Filter by control_standard

* Introduced in: 9.11
```

```
|pool
|string
|query
|False
a|Filter by pool


|aggregates.name
|string
|query
|False
a|Filter by aggregates.name


|aggregates.uuid
|string
|query
|False
a|Filter by aggregates.uuid


|outage.reason.arguments.message
|string
|query
|False
a|Filter by outage.reason.arguments.message

* Introduced in: 9.10


|outage.reason.arguments.code
|string
|query
|False
a|Filter by outage.reason.arguments.code

* Introduced in: 9.10


|outage.reason.code
|string
|query
|False
a|Filter by outage.reason.code
```

```
* Introduced in: 9.9


|outage.reason.message
|string
|query
|False
a|Filter by outage.reason.message

* Introduced in: 9.9


|outage.persistently_failed
|boolean
|query
|False
a|Filter by outage.persistently_failed

* Introduced in: 9.9


|state
|string
|query
|False
a|Filter by state


|name
|string
|query
|False
a|Filter by name


|storage_pool.name
|string
|query
|False
a|Filter by storage_pool.name

* Introduced in: 9.11


|storage_pool.uuid
|string
|query
```

```
|False
a|Filter by storage_pool.uuid

* Introduced in: 9.11


|vendor
|string
|query
|False
a|Filter by vendor


|shelf.uid
|string
|query
|False
a|Filter by shelf.uid


|right_size_sector_count
|integer
|query
|False
a|Filter by right_size_sector_count

* Introduced in: 9.11


|home_node.name
|string
|query
|False
a|Filter by home_node.name


|home_node.uuid
|string
|query
|False
a|Filter by home_node.uuid


|bytes_per_sector
|integer
|query
|False
```

```
a|Filter by bytes_per_sector

* Introduced in: 9.9


|stats.throughput
|integer
|query
|False
a|Filter by stats.throughput

* Introduced in: 9.9


|stats.path_error_count
|integer
|query
|False
a|Filter by stats.path_error_count

* Introduced in: 9.9


|stats.power_on_hours
|integer
|query
|False
a|Filter by stats.power_on_hours

* Introduced in: 9.9


|stats.iops_total
|integer
|query
|False
a|Filter by stats.iops_total

* Introduced in: 9.9


|stats.average_latency
|integer
|query
|False
a|Filter by stats.average_latency
```

```
* Introduced in: 9.9


|class
|string
|query
|False
a|Filter by class


|dr_node.name
|string
|query
|False
a|Filter by dr_node.name


|dr_node.uuid
|string
|query
|False
a|Filter by dr_node.uuid


|overall_security
|string
|query
|False
a|Filter by overall_security

* Introduced in: 9.11


|serial_number
|string
|query
|False
a|Filter by serial_number


|error.reason.arguments.message
|string
|query
|False
a|Filter by error.reason.arguments.message

* Introduced in: 9.10
```

```
|error.reason.arguments.code
|string
|query
|False
a|Filter by error.reason.arguments.code

* Introduced in: 9.10


|error.reason.code
|string
|query
|False
a|Filter by error.reason.code

* Introduced in: 9.9


|error.reason.message
|string
|query
|False
a|Filter by error.reason.message

* Introduced in: 9.9


|error.type
|string
|query
|False
a|Filter by error.type

* Introduced in: 9.9


|rpm
|integer
|query
|False
a|Filter by rpm


|firmware_version
|string
```

```
|query
|False
a|Filter by firmware_version


|self_encrypting
|boolean
|query
|False
a|Filter by self_encrypting

* Introduced in: 9.7


|key_id.fips
|string
|query
|False
a|Filter by key_id.fips

* Introduced in: 9.7


|key_id.data
|string
|query
|False
a|Filter by key_id.data

* Introduced in: 9.7


|container_type
|string
|query
|False
a|Filter by container_type


|usable_size
|integer
|query
|False
a|Filter by usable_size


|physical_size
```

```
|integer
|query
|False
a|Filter by physical_size

* Introduced in: 9.11


|model
|string
|query
|False
a|Filter by model


|fips_certified
|boolean
|query
|False
a|Filter by fips_certified

* Introduced in: 9.7


|drawer.id
|integer
|query
|False
a|Filter by drawer.id


|drawer.slot
|integer
|query
|False
a|Filter by drawer.slot


|compliance_standard
|string
|query
|False
a|Filter by compliance_standard

* Introduced in: 9.11
```

```
|type
|string
|query
|False
a|Filter by type


|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|location
|string
|query
|False
a|Filter by location

* Introduced in: 9.14


|effective_type
|string
|query
|False
a|Filter by effective_type

* Introduced in: 9.9


|paths.disk_path_name
|string
|query
|False
a|Filter by paths.disk_path_name

* Introduced in: 9.14
```

```
|paths.port_type
|string
|query
|False
a|Filter by paths.port_type

* Introduced in: 9.9


|paths.wwnn
|string
|query
|False
a|Filter by paths.wwnn

* Introduced in: 9.9


|paths.vmdisk_hypervisor_file_name
|string
|query
|False
a|Filter by paths.vmdisk_hypervisor_file_name

* Introduced in: 9.11


|paths.port_name
|string
|query
|False
a|Filter by paths.port_name

* Introduced in: 9.9


|paths.initiator
|string
|query
|False
a|Filter by paths.initiator

* Introduced in: 9.9


|paths.node.name
```

```
|string
|query
|False
a|Filter by paths.node.name

* Introduced in: 9.11


|paths.node.uuid
|string
|query
|False
a|Filter by paths.node.uuid

* Introduced in: 9.11


|paths.wwpn
|string
|query
|False
a|Filter by paths.wwpn

* Introduced in: 9.9


|uid
|string
|query
|False
a|Filter by uid


|virtual.storage_account
|string
|query
|False
a|Filter by virtual.storage_account

* Introduced in: 9.11


|virtual.target_address
|string
|query
|False
a|Filter by virtual.target_address
```

```
* Introduced in: 9.13


|virtual.object
|string
|query
|False
a|Filter by virtual.object

* Introduced in: 9.11


|virtual.container
|string
|query
|False
a|Filter by virtual.container

* Introduced in: 9.11


|local
|boolean
|query
|False
a|Filter by local

* Introduced in: 9.9


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
```

```
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
```

```
|integer
a|Number of records


|records
|array[link:#disk[disk]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "bay": 1,
      "bytes_per_sector": 520,
      "class": "solid_state",
      "compliance_standard": "FIPS 140-2",
      "container_type": "spare",
      "control_standard": "TCG Enterprise",
      "dr_node": {
        "name": "node1",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "effective_type": "vmdisk",
      "encryption_operation": "string",
      "error": [
        {
          "reason": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          },
          "type": "notallflashdisk"
        }
      ],
      "firmware_version": "NA51",
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "key_id": {
        "data": "string",
        "fips": "string"
      },
      "location": "node-01",
      "model": "X421_HCOBE450A10",
      "name": "1.0.1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "outage": {
```

```
        "reason": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        }
      },
      "overall_security": "Level 2",
      "paths": [
        {
          "disk_path_name": "vsim4:3a.10",
          "initiator": "3a",
          "node.name": "vsim4",
          "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
          "port_name": "A",
          "port_type": "sas",
          "vmdisk_hypervisor_file_name": "xvds vol0a0567ae156ca59f6",
          "wwnn": "5000c2971c1b2b8c",
          "wwpn": "5000c2971c1b2b8d"
        }
      ],
      "physical_size": 228930,
      "pool": "pool0",
      "protection_mode": "data",
      "rated_life_used_percent": 10,
      "right_size_sector_count": 1172123568,
      "rpm": 15000,
      "sector_count": 1172123568,
      "serial_number": "KHG2VX8R",
      "shelf": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uid": "7777841915827391056"
      },
      "state": "present",
      "stats": {
        "average_latency": 3,
        "iops_total": 12854,
        "path_error_count": 0,
```

```
        "power_on_hours": 21016,
        "throughput": 1957888
      },
      "storage_pool": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "storage_pool_1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "type": "ssd",
      "uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000",
      "usable_size": 959934889984,
      "vendor": "NETAPP",
      "virtual": {
        "container": "nviet12122018113936-rg",
        "object": "f1fu63se",
        "storage_account": "nviet12122018113936ps",
        "target_address": "string"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#dr_node]
[.api-collapsible-fifth-title]
dr_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
```

```
|integer
a|

|slot
|integer
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

The message and code detailing the error state of this disk.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#disk_error_info]
[.api-collapsible-fifth-title]
disk_error_info

[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|The message and code detailing the error state of this disk.


|type
|string
a|Disk error type.


|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#key_id]
[.api-collapsible-fifth-title]
key_id

[cols=3*,options=header]
|===
|Name
|Type
|Description

|data
|string
a|Key ID of the data authentication key


|fips
|string
a|Key ID of the FIPS authentication key


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error]
[.api-collapsible-fifth-title]
error

This error message and code explaining the disk failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|===


[#outage]
[.api-collapsible-fifth-title]
outage

Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|persistently_failed
|boolean
a|Indicates whether RAID maintains the state of this disk as failed across
reboots.


|reason
|link:#error[error]
a|This error message and code explaining the disk failure.


|===


[#disk_path_info]
[.api-collapsible-fifth-title]
disk_path_info

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_path_name
|string
a|Path-based disk name.


|initiator
```

```
|string
a|Initiator port.


|node.name
|string
a|Controller with the initiator port for this path.


|node.uuid
|string
a|Controller UUID, to identify node for this path.


|port_name
|string
a|Name of the disk port.


|port_type
|string
a|Disk port type.


|vmdisk_hypervisor_file_name
|string
a|Virtual disk hypervisor file name.


|wwnn
|string
a|Target device's World Wide Node Name.


|wwpn
|string
a|Target device's World Wide Port Name.


|===


[#shelf]
[.api-collapsible-fifth-title]
shelf

Shelf
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uid
|string
a|

|===


[#stats]
[.api-collapsible-fifth-title]
stats

[cols=3*,options=header]
|===
|Name
|Type
|Description

|average_latency
|integer
a|Average I/O latency across all active paths, in milliseconds.


|iops_total
|integer
a|Total I/O operations per second read and written to this disk across all
active paths.


|path_error_count
|integer
a|Disk path error count; failed I/O operations.


|power_on_hours
|integer
```

```
a|Hours powered on.


|throughput
|integer
a|Total disk throughput per second across all active paths, in bytes.


|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#virtual]
[.api-collapsible-fifth-title]
virtual

Information about backing storage for disks on cloud platforms.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|container
|string
a|Container name of the virtual disk.


|object
|string
a|Object name of the virtual disk.


|storage_account
|string
a|Storage account name of the virtual disk.


|target_address
|string
a|Target address of the virtual disk.


|===


[#disk]
[.api-collapsible-fifth-title]
disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of aggregates sharing this disk


|bay
|integer
a|Disk shelf bay
```

```
|bytes_per_sector
|integer
a|Bytes per sector.


|class
|string
a|Disk class


|compliance_standard
|string
a|Security standard that the device is certified to.


|container_type
|string
a|Type of overlying disk container


|control_standard
|string
a|Standard that the device supports for encryption control.


|dr_node
|link:#dr_node[dr_node]
a|

|drawer
|link:#drawer[drawer]
a|

|effective_type
|string
a|Effective Disk type


|encryption_operation
|string
a|This field should only be set as a query parameter in a PATCH operation.
It is input only and won't be returned by a subsequent GET.


|error
|array[link:#disk_error_info[disk_error_info]]
```

```
a|List of disk errors information.


|fips_certified
|boolean
a|

|firmware_version
|string
a|

|home_node
|link:#home_node[home_node]
a|

|key_id
|link:#key_id[key_id]
a|

|local
|boolean
a|Indicates if a disk is locally attached versus being remotely attached.
A locally attached disk resides in the same proximity as the host
cluster versus been attached to the remote cluster.


|location
|string
a|Physical location of the disk


|model
|string
a|

|name
|string
a|Cluster-wide disk name


|node
|link:#node[node]
a|

|outage
|link:#outage[outage]
a|Indicates if a disk has an entry in the failed disk registry, along with
```

```
the reason for the failure.


|overall_security
|string
a|Overall Security rating, for FIPS-certified devices.


|paths
|array[link:#disk_path_info[disk_path_info]]
a|List of paths to a disk


|physical_size
|integer
a|Physical size, in units of bytes


|pool
|string
a|Pool to which disk is assigned


|protection_mode
|string
a|Mode of drive data protection and FIPS compliance. Possible values are:

* _open_ - Data is unprotected
* _data_ - Data protection only, without FIPS compliance
* _part_ - Data is unprotected; other FIPS compliance settings present
* _full_ - Full data and FIPS compliance protection
* _miss_ - Protection mode information is not available


|rated_life_used_percent
|integer
a|Percentage of rated life used


|right_size_sector_count
|integer
a|Number of usable disk sectors that remain after subtracting the right-
size adjustment for this disk.


|rpm
|integer
```

```
a|Revolutions per minute


|sanitize_spare
|boolean
a|Confirms spare disk sanitization (non-cryptographically).


|sector_count
|integer
a|Number of sectors on the disk.


|self_encrypting
|boolean
a|

|serial_number
|string
a|

|shelf
|link:#shelf[shelf]
a|Shelf


|state
|string
a|State


|stats
|link:#stats[stats]
a|

|storage_pool
|link:#storage_pool[storage_pool]
a|Shared Storage Pool


|type
|string
a|Disk interface type


|uid
|string
```

```
a|The unique identifier for a disk


|usable_size
|integer
a|

|vendor
|string
a|

|virtual
|link:#virtual[virtual]
a|Information about backing storage for disks on cloud platforms.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID81f5be3e2f013d0eef4f3ed5a8717542]]
= Update disk ownership, change authentication keys, or sanitize disks

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/disks`#

*Introduced In:* 9.7

Updates disk ownership, changes authentication keys, sanitizes encrypting
disks, or sanitizes non-encrypting spare disks.

== Related ONTAP commands

* `storage disk assign`
* `storage disk removeowner`
* `storage encryption disk modify -data-key-id`
* `storage encryption disk sanitize`
* `security key-manager key query -key-type NSE-AK`
* `storage disk unfail`
* `storage disk sanitize-spare`

== Learn more

* xref:{relative_path}storage_disks_endpoint_overview.html[DOC
/storage/disks]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|name
|string
|query
|False
a|Disk name


|node
|string
|query
|False
a|Node to assign disk. This property is not supported on the ASA r2
platform.

* Introduced in: 9.8


|pool
|string
|query
|False
a|Pool to assign disk to

* Introduced in: 9.11


|encryption_operation
|string
|query
|False
a|Name of the operation to apply to encrypting disks.

* rekey_data_default changes the data authentication key (AK) to the
drive-unique Manufacture Secure ID (MSID) value. Allows the drive to be
attached to other clusters. Disables data-at-rest protection without
erasing the data.
* rekey_data_auto_id changes the data authentication key (AK) to an AK the
cluster selects automatically. Enables data-at-rest protection.
* sanitize_disk cryptographically erases all user data from a spare or
broken drive by altering the data encryption key. Resets the data AK to
the drive-unique MSID value and disables data-at-rest protection. Used
when a drive is being repurposed or returned.
* enum: ["rekey_data_default", "rekey_data_auto_id", "sanitize_disk"]


|sanitize_spare
```

```
|boolean
|query
|False
a|Confirms spare disk sanitization (non-cryptographically).

* Introduced in: 9.16


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of aggregates sharing this disk


|bay
|integer
a|Disk shelf bay


|bytes_per_sector
|integer
a|Bytes per sector.


|class
|string
a|Disk class
```

```
|compliance_standard
|string
a|Security standard that the device is certified to.


|container_type
|string
a|Type of overlying disk container


|control_standard
|string
a|Standard that the device supports for encryption control.


|dr_node
|link:#dr_node[dr_node]
a|

|drawer
|link:#drawer[drawer]
a|

|effective_type
|string
a|Effective Disk type


|encryption_operation
|string
a|This field should only be set as a query parameter in a PATCH operation.
It is input only and won't be returned by a subsequent GET.


|error
|array[link:#disk_error_info[disk_error_info]]
a|List of disk errors information.


|fips_certified
|boolean
a|

|firmware_version
|string
a|
```

```
|home_node
|link:#home_node[home_node]
a|

|key_id
|link:#key_id[key_id]
a|

|local
|boolean
a|Indicates if a disk is locally attached versus being remotely attached.
A locally attached disk resides in the same proximity as the host
cluster versus been attached to the remote cluster.


|location
|string
a|Physical location of the disk


|model
|string
a|

|name
|string
a|Cluster-wide disk name


|node
|link:#node[node]
a|

|outage
|link:#outage[outage]
a|Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


|overall_security
|string
a|Overall Security rating, for FIPS-certified devices.


|paths
|array[link:#disk_path_info[disk_path_info]]
```

```
a|List of paths to a disk


|physical_size
|integer
a|Physical size, in units of bytes


|pool
|string
a|Pool to which disk is assigned


|protection_mode
|string
a|Mode of drive data protection and FIPS compliance. Possible values are:

* _open_ - Data is unprotected
* _data_ - Data protection only, without FIPS compliance
* _part_ - Data is unprotected; other FIPS compliance settings present
* _full_ - Full data and FIPS compliance protection
* _miss_ - Protection mode information is not available


|rated_life_used_percent
|integer
a|Percentage of rated life used


|right_size_sector_count
|integer
a|Number of usable disk sectors that remain after subtracting the right-
size adjustment for this disk.


|rpm
|integer
a|Revolutions per minute


|sanitize_spare
|boolean
a|Confirms spare disk sanitization (non-cryptographically).


|sector_count
|integer
```

```
a|Number of sectors on the disk.


|self_encrypting
|boolean
a|

|serial_number
|string
a|

|shelf
|link:#shelf[shelf]
a|Shelf


|state
|string
a|State


|stats
|link:#stats[stats]
a|

|storage_pool
|link:#storage_pool[storage_pool]
a|Shared Storage Pool


|type
|string
a|Disk interface type


|uid
|string
a|The unique identifier for a disk


|usable_size
|integer
a|

|vendor
|string
a|
```

```
|virtual
|link:#virtual[virtual]
a|Information about backing storage for disks on cloud platforms.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "bay": 1,
  "bytes_per_sector": 520,
  "class": "solid_state",
  "compliance_standard": "FIPS 140-2",
  "container_type": "spare",
  "control_standard": "TCG Enterprise",
  "dr_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "effective_type": "vmdisk",
  "encryption_operation": "string",
  "error": [
    {
      "type": "notallflashdisk"
    }
  ],
  "firmware_version": "NA51",
  "home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "key_id": {
    "data": "string",
    "fips": "string"
  },
```

```
  "location": "node-01",
  "model": "X421_HCOBE450A10",
  "name": "1.0.1",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "overall_security": "Level 2",
  "paths": [
    {
      "disk_path_name": "vsim4:3a.10",
      "initiator": "3a",
      "node.name": "vsim4",
      "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
      "port_name": "A",
      "port_type": "sas",
      "vmdisk_hypervisor_file_name": "xvds vol0a0567ae156ca59f6",
      "wwnn": "5000c2971c1b2b8c",
      "wwpn": "5000c2971c1b2b8d"
    }
  ],
  "physical_size": 228930,
  "pool": "pool0",
  "protection_mode": "data",
  "rated_life_used_percent": 10,
  "right_size_sector_count": 1172123568,
  "rpm": 15000,
  "sector_count": 1172123568,
  "serial_number": "KHG2VX8R",
  "shelf": {
    "uid": "7777841915827391056"
  },
  "state": "present",
  "stats": {
    "average_latency": 3,
    "iops_total": 12854,
    "path_error_count": 0,
    "power_on_hours": 21016,
    "throughput": 1957888
  },
  "storage_pool": {
    "name": "storage_pool_1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "ssd",
  "uid":
```

```
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000",
  "usable_size": 959934889984,
  "vendor": "NETAPP",
  "virtual": {
    "container": "nviet12122018113936-rg",
    "object": "f1fu63se",
    "storage_account": "nviet12122018113936ps",
    "target_address": "string"
  }
}
====

== Response
```


Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 720951
| Unable to unfail the disk.

| 720968
| Quiet mode is applicable only when disk unfail is sparing the drive.

| 721029
| Failed to assign disks to a pool.

| 721066
| Node is outside the list of controllers for disk.

| 721067
| The specified node does not exist in the cluster.

| 721084
| An invalid pool name was specified.

| 721085
```

```
| The node name was not specified.

| 721093
| Operation not supported on this platform.

| 721095
| Failed to sanitize the spare disk (non-cryptographically).

| 721096
| Invalid disk state. This state is not supported for a PATCH operation.

| 1441795
| Setting the data key ID to the manufacture secure ID is not allowed when
in FIPS-compliance mode.

| 1441825
| The system does not support maintenance of the controls of encrypting
disks.

| 14155777
| The operation failed on one or more disks.

| 14155778
| No self-encrypting disks were specified.

| 14155779
| Status from a node shows that a conflicting operation has occurred. Some
disk controls might have changed.

| 14155780
| Could not retrieve the required key ID from the key manager.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#dr_node]
[.api-collapsible-fifth-title]
dr_node

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|slot
|integer
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

The message and code detailing the error state of this disk.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#disk_error_info]
[.api-collapsible-fifth-title]
disk_error_info
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|type
|string
a|Disk error type.


|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#key_id]
[.api-collapsible-fifth-title]
key_id

[cols=3*,options=header]
|===
|Name
|Type
|Description

|data
```

```
|string
a|Key ID of the data authentication key


|fips
|string
a|Key ID of the FIPS authentication key


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#error]
[.api-collapsible-fifth-title]
error

This error message and code explaining the disk failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#outage]
[.api-collapsible-fifth-title]
outage

Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|persistently_failed
|boolean
a|Indicates whether RAID maintains the state of this disk as failed across
reboots.


|===


[#disk_path_info]
[.api-collapsible-fifth-title]
disk_path_info

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|disk_path_name
|string
a|Path-based disk name.


|initiator
|string
a|Initiator port.


|node.name
|string
a|Controller with the initiator port for this path.


|node.uuid
|string
a|Controller UUID, to identify node for this path.


|port_name
|string
a|Name of the disk port.


|port_type
|string
a|Disk port type.


|vmdisk_hypervisor_file_name
|string
a|Virtual disk hypervisor file name.


|wwnn
|string
a|Target device's World Wide Node Name.


|wwpn
|string
a|Target device's World Wide Port Name.
```

```
|===


[#shelf]
[.api-collapsible-fifth-title]
shelf

Shelf


[cols=3*,options=header]
|===
|Name
|Type
|Description

|uid
|string
a|

|===


[#stats]
[.api-collapsible-fifth-title]
stats

[cols=3*,options=header]
|===
|Name
|Type
|Description

|average_latency
|integer
a|Average I/O latency across all active paths, in milliseconds.


|iops_total
|integer
a|Total I/O operations per second read and written to this disk across all
active paths.


|path_error_count
|integer
```

```
a|Disk path error count; failed I/O operations.


|power_on_hours
|integer
a|Hours powered on.


|throughput
|integer
a|Total disk throughput per second across all active paths, in bytes.


|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#virtual]
[.api-collapsible-fifth-title]
virtual

Information about backing storage for disks on cloud platforms.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|container
|string
a|Container name of the virtual disk.


|object
|string
a|Object name of the virtual disk.


|storage_account
|string
a|Storage account name of the virtual disk.


|target_address
|string
a|Target address of the virtual disk.


|===


[#disk]
[.api-collapsible-fifth-title]
disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of aggregates sharing this disk


|bay
|integer
a|Disk shelf bay
```

```
|bytes_per_sector
|integer
a|Bytes per sector.


|class
|string
a|Disk class


|compliance_standard
|string
a|Security standard that the device is certified to.


|container_type
|string
a|Type of overlying disk container


|control_standard
|string
a|Standard that the device supports for encryption control.


|dr_node
|link:#dr_node[dr_node]
a|

|drawer
|link:#drawer[drawer]
a|

|effective_type
|string
a|Effective Disk type


|encryption_operation
|string
a|This field should only be set as a query parameter in a PATCH operation.
It is input only and won't be returned by a subsequent GET.


|error
```

```
|array[link:#disk_error_info[disk_error_info]]
a|List of disk errors information.


|fips_certified
|boolean
a|

|firmware_version
|string
a|

|home_node
|link:#home_node[home_node]
a|

|key_id
|link:#key_id[key_id]
a|

|local
|boolean
a|Indicates if a disk is locally attached versus being remotely attached.
A locally attached disk resides in the same proximity as the host
cluster versus been attached to the remote cluster.


|location
|string
a|Physical location of the disk


|model
|string
a|

|name
|string
a|Cluster-wide disk name


|node
|link:#node[node]
a|

|outage
|link:#outage[outage]
```

```
a|Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


|overall_security
|string
a|Overall Security rating, for FIPS-certified devices.


|paths
|array[link:#disk_path_info[disk_path_info]]
a|List of paths to a disk


|physical_size
|integer
a|Physical size, in units of bytes


|pool
|string
a|Pool to which disk is assigned


|protection_mode
|string
a|Mode of drive data protection and FIPS compliance. Possible values are:

* _open_ - Data is unprotected
* _data_ - Data protection only, without FIPS compliance
* _part_ - Data is unprotected; other FIPS compliance settings present
* _full_ - Full data and FIPS compliance protection
* _miss_ - Protection mode information is not available


|rated_life_used_percent
|integer
a|Percentage of rated life used


|right_size_sector_count
|integer
a|Number of usable disk sectors that remain after subtracting the right-
size adjustment for this disk.


|rpm
```

```
|integer
a|Revolutions per minute


|sanitize_spare
|boolean
a|Confirms spare disk sanitization (non-cryptographically).


|sector_count
|integer
a|Number of sectors on the disk.


|self_encrypting
|boolean
a|

|serial_number
|string
a|

|shelf
|link:#shelf[shelf]
a|Shelf


|state
|string
a|State


|stats
|link:#stats[stats]
a|

|storage_pool
|link:#storage_pool[storage_pool]
a|Shared Storage Pool


|type
|string
a|Disk interface type


|uid
```

```
|string
a|The unique identifier for a disk


|usable_size
|integer
a|

|vendor
|string
a|

|virtual
|link:#virtual[virtual]
a|Information about backing storage for disks on cloud platforms.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
```

```
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID4e17b5e4c83e0b23ae410c80113ece5d]]
= Retrieve a specific disk

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/disks/{name}`#

*Introduced In:* 9.6

Retrieves a specific disk.

== Related ONTAP commands

* `storage disk show`
* `storage encryption disk show`

== Learn more

* xref:{relative_path}storage_disks_endpoint_overview.html[DOC
/storage/disks]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
```

```
a|Disk name


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of aggregates sharing this disk


|bay
|integer
a|Disk shelf bay


|bytes_per_sector
|integer
a|Bytes per sector.


|class
|string
a|Disk class


|compliance_standard
|string
a|Security standard that the device is certified to.
```

```
|container_type
|string
a|Type of overlying disk container


|control_standard
|string
a|Standard that the device supports for encryption control.


|dr_node
|link:#dr_node[dr_node]
a|

|drawer
|link:#drawer[drawer]
a|

|effective_type
|string
a|Effective Disk type


|encryption_operation
|string
a|This field should only be set as a query parameter in a PATCH operation.
It is input only and won't be returned by a subsequent GET.


|error
|array[link:#disk_error_info[disk_error_info]]
a|List of disk errors information.


|fips_certified
|boolean
a|

|firmware_version
|string
a|

|home_node
|link:#home_node[home_node]
a|

|key_id
```

```
|link:#key_id[key_id]
a|

|local
|boolean
a|Indicates if a disk is locally attached versus being remotely attached.
A locally attached disk resides in the same proximity as the host
cluster versus been attached to the remote cluster.


|location
|string
a|Physical location of the disk


|model
|string
a|

|name
|string
a|Cluster-wide disk name


|node
|link:#node[node]
a|

|outage
|link:#outage[outage]
a|Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


|overall_security
|string
a|Overall Security rating, for FIPS-certified devices.


|paths
|array[link:#disk_path_info[disk_path_info]]
a|List of paths to a disk


|physical_size
|integer
a|Physical size, in units of bytes
```

```
|pool
|string
a|Pool to which disk is assigned


|protection_mode
|string
a|Mode of drive data protection and FIPS compliance. Possible values are:

* _open_ - Data is unprotected
* _data_ - Data protection only, without FIPS compliance
* _part_ - Data is unprotected; other FIPS compliance settings present
* _full_ - Full data and FIPS compliance protection
* _miss_ - Protection mode information is not available


|rated_life_used_percent
|integer
a|Percentage of rated life used


|right_size_sector_count
|integer
a|Number of usable disk sectors that remain after subtracting the right-
size adjustment for this disk.


|rpm
|integer
a|Revolutions per minute


|sanitize_spare
|boolean
a|Confirms spare disk sanitization (non-cryptographically).


|sector_count
|integer
a|Number of sectors on the disk.


|self_encrypting
|boolean
a|
```

```
|serial_number
|string
a|

|shelf
|link:#shelf[shelf]
a|Shelf


|state
|string
a|State


|stats
|link:#stats[stats]
a|

|storage_pool
|link:#storage_pool[storage_pool]
a|Shared Storage Pool


|type
|string
a|Disk interface type


|uid
|string
a|The unique identifier for a disk


|usable_size
|integer
a|

|vendor
|string
a|

|virtual
|link:#virtual[virtual]
a|Information about backing storage for disks on cloud platforms.
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "bay": 1,
  "bytes_per_sector": 520,
  "class": "solid_state",
  "compliance_standard": "FIPS 140-2",
  "container_type": "spare",
  "control_standard": "TCG Enterprise",
  "dr_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "effective_type": "vmdisk",
  "encryption_operation": "string",
  "error": [
    {
      "reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "type": "notallflashdisk"
    }
  ],
```

```
  "firmware_version": "NA51",
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "key_id": {
    "data": "string",
    "fips": "string"
  },
  "location": "node-01",
  "model": "X421_HCOBE450A10",
  "name": "1.0.1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "outage": {
    "reason": {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist"
    }
  },
  "overall_security": "Level 2",
  "paths": [
    {
      "disk_path_name": "vsim4:3a.10",
      "initiator": "3a",
      "node.name": "vsim4",
      "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
      "port_name": "A",
```

```
      "port_type": "sas",
      "vmdisk_hypervisor_file_name": "xvds vol0a0567ae156ca59f6",
      "wwnn": "5000c2971c1b2b8c",
      "wwpn": "5000c2971c1b2b8d"
    }
  ],
  "physical_size": 228930,
  "pool": "pool0",
  "protection_mode": "data",
  "rated_life_used_percent": 10,
  "right_size_sector_count": 1172123568,
  "rpm": 15000,
  "sector_count": 1172123568,
  "serial_number": "KHG2VX8R",
  "shelf": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uid": "7777841915827391056"
  },
  "state": "present",
  "stats": {
    "average_latency": 3,
    "iops_total": 12854,
    "path_error_count": 0,
    "power_on_hours": 21016,
    "throughput": 1957888
  },
  "storage_pool": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "storage_pool_1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "ssd",
  "uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000",
  "usable_size": 959934889984,
  "vendor": "NETAPP",
  "virtual": {
```

```
    "container": "nviet12122018113936-rg",
    "object": "f1fu63se",
    "storage_account": "nviet12122018113936ps",
    "target_address": "string"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#dr_node]
[.api-collapsible-fifth-title]
dr_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|id
|integer
a|

|slot
|integer
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

The message and code detailing the error state of this disk.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#disk_error_info]
[.api-collapsible-fifth-title]
disk_error_info

[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|The message and code detailing the error state of this disk.


|type
|string
a|Disk error type.


|===


[#home_node]
[.api-collapsible-fifth-title]
```

```
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#key_id]
[.api-collapsible-fifth-title]
key_id

[cols=3*,options=header]
|===
|Name
|Type
|Description

|data
|string
a|Key ID of the data authentication key


|fips
|string
a|Key ID of the FIPS authentication key


|===


[#node]
```

```
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error]
[.api-collapsible-fifth-title]
error

This error message and code explaining the disk failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|===


[#outage]
[.api-collapsible-fifth-title]
outage

Indicates if a disk has an entry in the failed disk registry, along with
the reason for the failure.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|persistently_failed
|boolean
a|Indicates whether RAID maintains the state of this disk as failed across
reboots.


|reason
|link:#error[error]
a|This error message and code explaining the disk failure.


|===


[#disk_path_info]
[.api-collapsible-fifth-title]
disk_path_info

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_path_name
```

```
|string
a|Path-based disk name.


|initiator
|string
a|Initiator port.


|node.name
|string
a|Controller with the initiator port for this path.


|node.uuid
|string
a|Controller UUID, to identify node for this path.


|port_name
|string
a|Name of the disk port.


|port_type
|string
a|Disk port type.


|vmdisk_hypervisor_file_name
|string
a|Virtual disk hypervisor file name.


|wwnn
|string
a|Target device's World Wide Node Name.


|wwpn
|string
a|Target device's World Wide Port Name.


|===
```

```
[#shelf]
[.api-collapsible-fifth-title]
shelf

Shelf


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uid
|string
a|

|===


[#stats]
[.api-collapsible-fifth-title]
stats

[cols=3*,options=header]
|===
|Name
|Type
|Description

|average_latency
|integer
a|Average I/O latency across all active paths, in milliseconds.


|iops_total
|integer
a|Total I/O operations per second read and written to this disk across all
active paths.


|path_error_count
|integer
```

```
a|Disk path error count; failed I/O operations.


|power_on_hours
|integer
a|Hours powered on.


|throughput
|integer
a|Total disk throughput per second across all active paths, in bytes.


|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

Shared Storage Pool


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#virtual]
[.api-collapsible-fifth-title]
virtual
```

```
Information about backing storage for disks on cloud platforms.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|container
|string
a|Container name of the virtual disk.


|object
|string
a|Object name of the virtual disk.


|storage_account
|string
a|Storage account name of the virtual disk.


|target_address
|string
a|Target address of the virtual disk.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage file clones

:leveloffset: +1


[[IDf54b5ebd2eb153eb8b2e038dc0f9463e]]
= Manage storage file clones



== Overview

You can use these API's to create file clones, retrieve split status and
manage split loads. These endpoints are used for cloning files within a
volume, without taking much of extra space. Child and parent clones shares
the unchanged blocks of data. REST support for sis clone creation for Flex
Group volumes is blocked.

A file clone split operation detach child clone from its parent. Split
operations use space. To ensure that file clone create operation is not
```

```
affected by split, file clone tokens are use to reserve space. API
endpoints can be used to update the validity and space reserved by token.

== File clone APIs

The following APIs are used to perform the following operations:

&ndash; POST      /api/storage/file/clone

&ndash;  GET      /api/storage/file/clone/split-status

&ndash; PATCH     /api/storage/file/clone/split-loads/{node.uuid}

&ndash;  GET      /api/storage/file/clone/split-loads/{node.uuid}

&ndash;  GET      /api/storage/file/clone/split-loads

&ndash;  GET      /api/storage/file/clone/tokens/

&ndash; DELETE    /api/storage/file/clone/tokens/{node.uuid}/{token.uuid}

&ndash; PATCH     /api/storage/file/clone/tokens/{node.uuid}/{token.uuid}




[[ID583e3675f2cc9ad1b3b5e03baf8a3030]]
= Create a clone of the file

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/file/clone`#

*Introduced In:* 9.6

Creates a clone of the file.

== Required Properties

* `source_path`
* `destination_path`
* `volume.uuid` - Instance UUID of the volume in which to create clone.

== Optional Properties

* `range` -  Required only in the case of a sub file clone.
* `autodelete` - Marks a cloned file for auto deletion.
```

```
* `is_backup` - Cloned file is used as a backup.
* `volume.name` - Name of the volume in which to create the clone.
* `token_uuid` - UUID of existing clone token.

== Related Ontap commands

* `volume file clone create`

== Creating file clones

The POST operation is used to create file clones with the specified
attributes in the body. Set the `volume.uuid` to identify the volume.
Set `source_path` and `destination_path` to identify the file path of the
original and copied file. For a full file clone, the new file is created
using `destination_path`.
For a sub file clone, set `range` in the format source-file-block-
number:destination-file-block-number:block-count. The API returns an error
for the following overlapping conditions: (a) if the source and
destination files are the same and if any of the source ranges  overlap
with any of the destination ranges. (b) if any of the source ranges
overlap amongst themselves. (c) if any of the destination ranges overlap
amongst themselves. If not provided, full file cloning is assumed.

If set to `autodelete`, the cloned file is deleted when the volumes are
full.<br>

----

# The API:
curl -X POST "https://<mgmt_ip>/api/storage/file/clone" -H "accept:
application/hal+json" -d '{"volume": {"name": "vol1",  "uuid": "40e0fdc5-
c28f-11eb-8270-005056bbeb0b"}, "source_path": "f1", "destination_path":
"f2_c1"}'

# The response:
{
"job": {
  "uuid": "0d025fd9-c4dc-11eb-adb5-005056bbeb0b",
  "_links": {
     "self": {
       "href": "/api/cluster/jobs/0d025fd9-c4dc-11eb-adb5-005056bbeb0b"
     }
  }
}
}
----
```

```
== Learn More

* [`DOC /storage/file/clone`]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|autodelete
|boolean
a|Mark clone file for auto deletion.


|destination_path
|string
a|Relative path of the clone/destination file in the volume.


|is_backup
|boolean
a|Mark clone file for backup.


|overwrite_destination
|boolean
a|Destination file gets overwritten.


|range
|array[string]
a|List of block ranges for sub-file cloning in the format "source-file-
block-number:destination-file-block-number:block-count"


|source_path
|string
a|Relative path of the source file in the volume.


|token_uuid
|string
a|UUID of existing clone token with reserved split load.


|volume
|link:#volume[volume]
a|

|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "destination_path": "dest_file1 or dir1/dest_file2",
  "is_backup": "",
  "range": [
    2
  ],
  "source_path": "src_file1 or dir1/src_file2 or
./.snapshot/snap1/src_file3",
  "token_uuid": "string",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 917812
| A relative path must be used for `destination_path` when a volume is
specified.

| 917864
| Failed to get the file handle.

| 917898
| Invalid `value` for the field -range. Value must be provided in the
format <source start block>:<destination start block>:<block length>.

| 918236
| The specified `volume.uuid` and `volume.name` refer to different
volumes.

| 13565952
| Clone start failed.

| 13565988
| The destination file cannot be marked as a backup clone as only a full-
file clone can be a backup clone.

| 13565998
| File clone using autodelete is not supported on a FlexGroup volume. Try
the operation again without specifying the "autodelete" parameter.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
* x-nullable: true


|===


[#file_clone]
[.api-collapsible-fifth-title]
file_clone

File clone


[cols=3*,options=header]
|===
|Name
|Type
|Description

|autodelete
|boolean
a|Mark clone file for auto deletion.


|destination_path
|string
a|Relative path of the clone/destination file in the volume.


|is_backup
|boolean
a|Mark clone file for backup.


|overwrite_destination
|boolean
a|Destination file gets overwritten.


|range
|array[string]
a|List of block ranges for sub-file cloning in the format "source-file-
block-number:destination-file-block-number:block-count"


|source_path
|string
```

```
a|Relative path of the source file in the volume.


|token_uuid
|string
a|UUID of existing clone token with reserved split load.


|volume
|link:#volume[volume]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDad405eddeea55bd3dbaf77d52220512e]]
= Retrieve the clone split load of a node

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/split-loads`#

*Introduced In:* 9.10

Retrieves the clone split load of a node.

== Related Ontap Commands

* `volume file clone split load show`

== Retrieving file clone split load related information

The GET operation can be used to retrieve information about clone split
load data. Split load data is the data currently undergoing the split.
There is a limit on split load data. This API communicates how much data
is undergoing split and how much can still be processed.

----

# The API:
/api/storage/file/clone/split-loads

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-loads" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "158d592f-a829-11eb-a47b-005056bb46d7",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/158d592f-a829-11eb-a47b-
005056bb46d7"
        }
      }
    },
    "load": {
      "maximum": 35184372088832,
```

```
      "current": 0,
      "token_reserved": 0,
      "allowable": 35184372088832
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-loads/158d592f-a829-11eb-
a47b-005056bb46d7"
      }
    }
  },
  {
    "node": {
      "uuid": "9686b8d1-a828-11eb-80d8-005056bbe7b6",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/9686b8d1-a828-11eb-80d8-
005056bbe7b6"
        }
      }
    },
    "load": {
      "maximum": 35184372088832,
      "current": 0,
      "token_reserved": 0,
      "allowable": 35184372088832
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-loads/9686b8d1-a828-11eb-
80d8-005056bbe7b6"
      }
    }
  }
],
"num_records": 2,
"_links":
  "self": {
    "href": "/api/storage/file/clone/split-loads"
  }
}
}
----

== Learn More
```

```
* [`DOC /storage/file/clone`]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|load.allowable
|integer
|query
|False
a|Filter by load.allowable


|load.token_reserved
|integer
|query
|False
a|Filter by load.token_reserved


|load.maximum
|integer
|query
|False
a|Filter by load.maximum


|load.current
|integer
|query
|False
a|Filter by load.current


|node.name
|string
|query
|False
```

```
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#split_load[split_load]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "load": {
        "allowable": "4KB",
        "current": 0,
        "maximum": "4KB",
        "token_reserved": 0
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#load]
[.api-collapsible-fifth-title]
load

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|allowable
|integer
a|Specifies the available file clone split load on the node.


|current
|integer
a|Specifies the current on-going file clone split load on the node.


|maximum
|integer
a|Specifies the maximum allowable file clone split load on the node at any
point in time. The least allowable file clone split load is 4KB and the
maximum is 675TB. The default value is 4KB for the maximum split load when
unit is not specified. The default file clone split load is set based on
the system configuration.


|token_reserved
|integer
a|Specifies the file clone split load on the node reserved for tokens.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
```

```
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#split_load]
[.api-collapsible-fifth-title]
split_load

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|load
|link:#load[load]
a|

|node
|link:#node[node]
a|

|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID87153e5f009d84ae778942680c2b9911]]
= Retrieve the volume file clone split load

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/split-loads/{node.uuid}`#

*Introduced In:* 9.10

Retrieve the volume file clone split load.

== Related ONTAP command

* `volume file clone split load show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|load
|link:#load[load]
a|

|node
|link:#node[node]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "load": {
    "allowable": "4KB",
```

```
    "current": 0,
    "maximum": "4KB",
    "token_reserved": 0
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
```

```
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#load]
[.api-collapsible-fifth-title]
load

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allowable
|integer
a|Specifies the available file clone split load on the node.


|current
|integer
a|Specifies the current on-going file clone split load on the node.


|maximum
|integer
a|Specifies the maximum allowable file clone split load on the node at any
point in time. The least allowable file clone split load is 4KB and the
maximum is 675TB. The default value is 4KB for the maximum split load when
unit is not specified. The default file clone split load is set based on
the system configuration.


|token_reserved
|integer
a|Specifies the file clone split load on the node reserved for tokens.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID0301fbfedad9e587e2cc739c9b51b4a0]]
= Update the maximum split load

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/file/clone/split-loads/{node.uuid}`#

*Introduced In:* 9.10

Updates the maximum split load.

== Related Ontap command

* `volume file clone split load modify`

== Learn More

* [`DOC /storage/file/clone`]

----

# The call:
curl -X PATCH "https://<mgmt_IP>/api/storage/file/clone/split-
loads/9686b8d1-a828-11eb-80d8-005056bbe7b6" -d '{"load": {"maximum":
"16TB" } }'

# The response to successful patch is empty body
----


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|load
|link:#load[load]
a|

|node
|link:#node[node]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "load": {
    "allowable": "4KB",
    "current": 0,
    "maximum": "4KB",
    "token_reserved": 0
  },
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13565966
| Failed to modify the maximum split load on the node due to an invalid
field value.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#load]
[.api-collapsible-fifth-title]
```

```
load

[cols=3*,options=header]
|===
|Name
|Type
|Description

|allowable
|integer
a|Specifies the available file clone split load on the node.


|current
|integer
a|Specifies the current on-going file clone split load on the node.


|maximum
|integer
a|Specifies the maximum allowable file clone split load on the node at any
point in time. The least allowable file clone split load is 4KB and the
maximum is 675TB. The default value is 4KB for the maximum split load when
unit is not specified. The default file clone split load is set based on
the system configuration.


|token_reserved
|integer
a|Specifies the file clone split load on the node reserved for tokens.


|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|

|uuid
|string
a|

|===


[#split_load]
[.api-collapsible-fifth-title]
split_load

[cols=3*,options=header]
|===
|Name
|Type
|Description

|load
|link:#load[load]
a|

|node
|link:#node[node]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID8fa9937c539944d97cb371e76bee319f]]
= Retrieve file clone split status of all volumes in the node

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/split-status`#

*Introduced In:* 9.10

Retrieves file clone split status of all volumes in the node.

== Learn More

* [`DOC /storage/file/clone`]

----

# The API:
/api/storage/file/clone/split-status

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-status" -H
"accept: application/hal+json"

# The response:
{
"records": [
  {
    "volume": {
      "uuid": "ac559964-57a3-40cf-b5cb-f3cb99151a7d",
      "name": "vol1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/ac559964-57a3-40cf-b5cb-
f3cb99151a7d"
        }
      }
    },
    "svm": {
      "uuid": "9fcb44bf-4305-11e8-b8d5-00a09887594b",
      "name": "vs1"
      "_links": {
        "self": {
          "href": "/api/svm/svms/9fcb44bf-4305-11e8-b8d5-00a09887594b"
        }
      }
    },
```

```
    "pending_splits": 0,
    "unsplit_size": 0,
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-status/ac559964-57a3-40cf-
b5cb-f3cb99151a7d"
      }
    }
  },
  {
    "volume": {
      "uuid": "32d95d48-d8b7-11eb-a41d-005056bb3837",
      "name": "vs1_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/32d95d48-d8b7-11eb-a41d-
005056bb3837"
        }
      }
    },
    "svm": {
      "name": "vs1"
      "_links": {
        "self": {
          "href": "/api/svm/svms/9fcb44bf-4305-11e8-b8d5-00a09887594b"
        }
      }
    },
    "pending_splits": 0,
    "unsplit_size": 0,
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-status/32d95d48-d8b7-11eb-
a41d-005056bb3837"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/file/clone/split-status"
  }
}
}
----
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|volume.name
|string
|query
|False
a|Filter by volume.name


|volume.uuid
|string
|query
|False
a|Filter by volume.uuid


|unsplit_size
|integer
|query
|False
a|Filter by unsplit_size
```

```
* Introduced in: 9.15


|pending_splits
|integer
|query
|False
a|Filter by pending_splits


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records


|records
|array[link:#split_status[split_status]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "pending_splits": 0,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unsplit_size": 0,
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#split_status]
[.api-collapsible-fifth-title]
split_status

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|pending_splits
|integer
```

```
a|Specifies the number of pending file clone split operations in the
volume.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|unsplit_size
|integer
a|Specifies the space occupied by unsplit file clones in the volume.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf877091793f76fd7870ee024989337ee]]
= Retrieve file clone split status of all volumes

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/split-status/{volume.uuid}`#

*Introduced In:* 9.10

Retrieves file clone split status of all volumes in the node.

== Related Ontap commands
```

```
* `volume file clone split status`

== Learn More

* [`DOC /storage/file/clone`]

== Retrieves the information of split status.

The GET operation retrieves information about split processes in the
volume.
`pending-clone-splits` is the number of files for which file clone split
is not yet completed.
`unsplit-size` is the sum of all sizes in the volume that is not split, in
bytes.

----

# The API:
/api/storage/file/clone/split-status/{volume.uuid}

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-
status/ac559964-57a3-40cf-b5cb-f3cb99151a7d" -H "accept:
application/hal+json"

# The response:
{
"volume": {
  "uuid": "ac559964-57a3-40cf-b5cb-f3cb99151a7d",
  "name": "vol1",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/ac559964-57a3-40cf-b5cb-f3cb99151a7d"
    }
  }
},
"svm": {
  "name": "vs1"
},
"pending_splits": 0,
"unsplit_size": 0,
"_links": {
  "self": {
    "href": "/api/storage/file/clone/split-status/ac559964-57a3-40cf-b5cb-
f3cb99151a7d"
  }
```

```
}
}
----


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume Instance UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|pending_splits
```

```
|integer
a|Specifies the number of pending file clone split operations in the
volume.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|unsplit_size
|integer
a|Specifies the space occupied by unsplit file clones in the volume.


|volume
|link:#volume[volume]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "pending_splits": 0,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unsplit_size": 0,
  "volume": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID3c5c699abc9f640ffb541ee20e878b14]]
= Retrieve information for a token

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/tokens`#
```

```
*Introduced In:* 9.10

Retrieves information for the specified token.

== Related Ontap command

* `volume file clone token show`

== Learn More

* [`DOC /storage/file/clone`]

== Retrieving information on clone tokens

----

# The API:
/api/storage/file/clone/tokens

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/tokens" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "97255711-a1ad-11eb-92b2-0050568eb2ca",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/97255711-a1ad-11eb-92b2-
0050568eb2ca"
        }
      }
    },
    "uuid": "905c42ce-a74b-11eb-bd86-0050568ec7ae",
    "reserve_size": 10240,
    "expiry_time": {
      "limit": "PT1H10M",
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"
```

```
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/file/clone/tokens"
  }
}
}
```


```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|reserve_size
|integer
|query
|False
a|Filter by reserve_size


|uuid
|string
|query
|False
a|Filter by uuid


|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
```

```
|string
|query
|False
a|Filter by node.uuid


|expiry_time.left
|string
|query
|False
a|Filter by expiry_time.left


|expiry_time.limit
|string
|query
|False
a|Filter by expiry_time.limit


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
```

```
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records.


|records
|array[link:#token[token]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {}
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#expiry_time]
[.api-collapsible-fifth-title]
expiry_time

[cols=3*,options=header]
|===
|Name
|Type
|Description

|left
|string
a|Specifies the time remaining before the given token expires in ISO-8601
format.


|limit
|string
a|Specifies when the given token expires in ISO-8601 format.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Node name


|uuid
|string
a|Node UUID


|===
```

```
[#token]
[.api-collapsible-fifth-title]
token

token


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|expiry_time
|link:#expiry_time[expiry_time]
a|

|node
|link:#node[node]
a|

|reserve_size
|integer
a|Specifies the available reserve in the file clone split load for the
given token. Also note that the minimum value for reserve size is 4KB and
any value specified below 4KB will be rounded off to 4KB.


|uuid
|string
a|Token UUID.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID9fa83ac4a694cc7b7d83b21a431f06b2]]
= Create a new token to reserve the split load

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/file/clone/tokens`#

*Introduced In:* 9.10

Creates a new token to reserve the split load.

== Required Properties

* `node.uuid`
* `reserve-size`

== Optional Properties

* `expiry_time.limit`

== Default values

* `expiry_time.limit` - "60"

== Related ONTAP Commands

* `volume file clone token create`

== Learn More

* [`DOC /storage/file/clone`]

== Creating clone tokens to reserve space for clone creation on the node

There is a limit on the amount of clone data that can undergo a split at a
point of time on the node (clone split load). Clone tokens are used to
reserve space from clone split load for clone creation. The POST operation
is used to create clone tokens with `reserve-size` and `expiry-time.limit`
in the body.<br>
```

```
----

# The API
/api/storage/file/clone/tokens

# The call
curl -X POST "https://<mgmt_ip>/api/storage/file/clone/tokens" -H "accept:
application/hal+json" -d '{"node": {"uuid": "97255711-a1ad-11eb-92b2-
0050568eb2ca"}, "reserve_size": "40M", "expiry_time": { "limit": "4200"}
}'

# The response
{
"num_records": 1,
"records": [
  {
    "node": {
      "uuid": "97255711-a1ad-11eb-92b2-0050568eb2ca",
      "name": "node1"
    },
    "uuid": "286f6ae4-c94d-11eb-adb5-005056bbeb0b",
    "reserve_size": 41943040,
    "expiry_time": {
      "limit": "PT1H10M"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca"
      }
    }
  }
]
}
----


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
```

```
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|link:#expiry_time[expiry_time]
a|

|node
|link:#node[node]
a|

|reserve_size
|integer
a|Specifies the available reserve in the file clone split load for the
given token. Also note that the minimum value for reserve size is 4KB and
any value specified below 4KB will be rounded off to 4KB.


|uuid
|string
a|Token UUID.


|===


== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.


|records
|array[link:#token[token]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "num_records": 1,
  "records": [
    {}
  ]
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
```

```
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13565959
| Failed to create token for node `node.name`.

| 13565964
| Reserve size cannot be negative.

| 13565997
| The node.uuid doesn't exist.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
[#expiry_time]
[.api-collapsible-fifth-title]
expiry_time

[cols=3*,options=header]
|===
|Name
|Type
|Description

|left
|string
a|Specifies the time remaining before the given token expires in ISO-8601
format.


|limit
|string
a|Specifies when the given token expires in ISO-8601 format.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Node name
```

```
|uuid
|string
a|Node UUID


|===


[#token]
[.api-collapsible-fifth-title]
token

token


[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|link:#expiry_time[expiry_time]
a|

|node
|link:#node[node]
a|

|reserve_size
|integer
a|Specifies the available reserve in the file clone split load for the
given token. Also note that the minimum value for reserve size is 4KB and
any value specified below 4KB will be rounded off to 4KB.


|uuid
|string
a|Token UUID.


|===


[#collection_links]
[.api-collapsible-fifth-title]
```

```
collection_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID2fde745a61734a4f9677b77aa297145e]]
= Delete a specific file clone token

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/file/clone/tokens/{node.uuid}/{uuid}`#

*Introduced In:* 9.10

Deletes a specific file clone token.

== Related Ontap command

* `volume file clone token delete`

== Delete specific clone token.

----

# The API:
/api/storage/file/clone/tokens/{node.uuid}/{token.uuid}

# The call:
curl -X DELETE "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
a1ad-11eb-92b2-0050568eb2ca/909c42ce-a74b-11eb-bd86-0050568ec7ae"

# The successful response is empty body.
----

== Learn More

* [`DOC /storage/file/clone`]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|uuid
|string
|path
|True
a|Token UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13565958
| Failed to get information about token `uuid` for node `node.name`.

| 13565961
| Failed to delete token for node `node.name`.
```

```
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID62e48083bcd9931267593df745b90c0a]]
= Retrieve a file clone token

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/clone/tokens/{node.uuid}/{uuid}`#

*Introduced In:* 9.10

Retrieves a file clone token

== Related Ontap command

* `volume file clone token show`

== Retrieve information for single token.

----

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
a1ad-11eb-92b2-0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"

# The response:
{
"node": {
  "uuid": "97255711-a1ad-11eb-92b2-0050568eb2ca",
  "name": "node1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/97255711-a1ad-11eb-92b2-0050568eb2ca"
    }
  }
},
```

```
"uuid": "905c42ce-a74b-11eb-bd86-0050568ec7ae",
"reserve_size": 41943040,
"expiry_time": {
  "limit": "PT1H10M",
},
"_links": {
  "self": {
    "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"
  }
}
}
```


```
----

== Learn More

* [`DOC /storage/file/clone`]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|uuid
|string
|path
|True
a|Token UUID


|fields
|array[string]
|query
```

```
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|expiry_time
|link:#expiry_time[expiry_time]
a|

|node
|link:#node[node]
a|

|reserve_size
|integer
a|Specifies the available reserve in the file clone split load for the
given token. Also note that the minimum value for reserve size is 4KB and
any value specified below 4KB will be rounded off to 4KB.


|uuid
|string
a|Token UUID.


|===


== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13565958
| Failed to get information about token `uuid` for node `node.name`.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#expiry_time]
[.api-collapsible-fifth-title]
expiry_time

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|left
|string
a|Specifies the time remaining before the given token expires in ISO-8601
format.


|limit
|string
a|Specifies when the given token expires in ISO-8601 format.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Node name


|uuid
|string
a|Node UUID


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID3956cee01d936794d38fecad5ed49523]]
= Update a file clone token

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/file/clone/tokens/{node.uuid}/{uuid}`#

*Introduced In:* 9.10

Updates a file clone token.

== Related Ontap commands

* `volume file clone token modify`

== Modify clone token

Use the PATCH API to update the expiry time associated with the clone
token.<br>

----

# The call:
curl -X PATCH "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
a1ad-11eb-92b2-0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae" -d
'{"expiry_time": {"limit": "5400"} }'

# The response for successful PATCH is empty.
----

== Learn More

* [`DOC /storage/file/clone`]


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|uuid
|string
|path
|True
a|Token UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|link:#expiry_time[expiry_time]
a|

|uuid
|string
a|Token UUID.


|===


== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 13565958
| Failed to get information about token `uuid` for node `node.name`.

| 13565960
| Failed to modify token expiry limit for node `node.name`.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#self_link]
[.api-collapsible-fifth-title]
self_link
[#expiry_time]
[.api-collapsible-fifth-title]
expiry_time

[cols=3*,options=header]
|===
|Name
|Type
|Description

|left
|string
a|Specifies the time remaining before the given token expires in ISO-8601
format.


|limit
|string
a|Specifies when the given token expires in ISO-8601 format.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Node name


|uuid
|string
a|Node UUID
```

```
|===


[#token]
[.api-collapsible-fifth-title]
token

token


[cols=3*,options=header]
|===
|Name
|Type
|Description

|expiry_time
|link:#expiry_time[expiry_time]
a|

|uuid
|string
a|Token UUID.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
```

```
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1
```

```
[[ID10afafeb60534e66d496b96b43d96e5f]]
= Start a file copy operation

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/file/copy`#

*Introduced In:* 9.8

Starts a file copy operation. Only supported on flexible volumes.

== Required properties

* `files_to_copy` - List of files with the destination they are to be
copied to.

== Default property values

* `cutover_time` - _10_
* `hold_quiescence` - _false_
* `max_throughput` - _0_
* `reference_cutover_time` - _10_

== Related ONTAP commands

* `volume file copy start`

== Examples

== Copying two files

The POST request is used to copy file(s).

----

# The API:
/api/storage/file/copy

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/copy" -H "accept:
application/hal+json" -d
'{"files_to_copy":[{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0
"},"path":"d1/src_f1"},"destination":{"volume":{"name":"vol_a"},"svm":{"na
me":"vs0"},"path":"d1/dst_f1"}},
{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"d1/src_f
2"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"
d1/dst_f2"}}]}'
```

```
# The response:
{
"job": {
  "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
  "_links": {
     "self": {
       "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
     }
   }
 }
}
----


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
```

```
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|The maximum amount of time (in seconds) that the source can be quiesced
before a destination file must be made available for read-write traffic.


|files_to_copy
|array[link:#files_to_copy[files_to_copy]]
a|A list of source files along with the destinations they are copied to.
If the terminal path component of the destination is a directory, then the
source file's basename is replicated in that directory.


|hold_quiescence
|boolean
a|Specifies whether the source file should be held quiescent for the
duration of the copy operation.


|max_throughput
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1MB/s is set to
1MB/s. A non-zero value greater than 1MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.


|reference_cutover_time
|integer
a|The maximum amount of time (in seconds) that the source reference file
```

```
can be quiesced before the corresponding destination file must be made
available for read-write traffic.


|reference_path
|string
a|The source reference file. If a reference file is specified, data for
other files being copied will be transferred as a difference from the
reference file. This can save bandwidth and destination storage if the
specified source files share blocks. If provided, this input must match
one of the source file paths. This input need not be provided if only one
source file is specified.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "cutover_time": 10,
  "files_to_copy": [
    {
      "destination": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "source": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
```

```
      }
    }
  ],
  "reference_cutover_time": 10,
  "reference_path": "svm1:volume1/file1"
}
====
```

```
== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 917506
| Volume `volume.name` not found.

| 918236
| The specified `volume.uuid` and `volume.name` refer to different
volumes.

| 2621464
| Failed to lookup identifier for SVM `svm.name`.

| 2621706
| The specified `svm.uuid` and `svm.name` refer to different SVMs.

| 7012352
| File locations are inconsistent. All files must be on the same volume.

| 7012353
| Exceeded the file operations supported number of files.

| 7012354
| Unable to pair the number of source files to destination files.

| 7012357
| Cannot start a file operation until all cluster nodes support the file
operations capability.

| 7012358
| The specified source path is invalid.

| 7012359
| The specified destination path is invalid.

| 7012362
| The SVMs are not yet in a peered state yet.

| 7012363
```

```
| Cannot copy files. All file operations must be managed by the
destination SVM's administrator.

| 7012365
| Copying a file between clusters is not supported.

| 7012366
| Cannot start the operation with an invalid "max-throughput" value.

| 7012367
| A reference path may only be specified if multiple source paths are
specified.

| 7012368
| The reference path must have a matching source path.

| 7012371
| The reference cutover time exceeds the maximum allowable time.

| 7012374
| Source volume and destination volume have different home clusters.

| 7012376
| Operation not allowed on a volume that is part of a SnapMirror
Synchronous relationship.

| 7012377
| Cannot start a file copy operation on the volume because an active
volume conversion is in progress.

| 7012378
| The source volume UUID or name must be provided.

| 7012380
| The source SVM UUID or name must be provided.

| 7013360
| The peering relationship between the two SVMs does not include
application "file-copy".

| 7013361
| The SVMs are not in an intracluster peering relationship.

| 7018866
| Cannot start the file operation. The volume is read-only.

| 7018877
```

```
| Maximum combined total (50) of file and LUN copy and move operations
reached. When one or more of the operations has completed, try the command
again.

| 13107223
| Operation not supported for FlexGroup volumes or FlexGroup constituents.

| 196608143
| Cannot start operation. The volume is undergoing a secure purge
operation.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
* x-nullable: true


|===


[#file_reference]
[.api-collapsible-fifth-title]
file_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|Path of the file or directory.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#files_to_copy]
[.api-collapsible-fifth-title]
files_to_copy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|link:#file_reference[file_reference]
a|
```

```
|source
|link:#file_reference[file_reference]
a|

|===


[#file_copy]
[.api-collapsible-fifth-title]
file_copy

File copy


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|The maximum amount of time (in seconds) that the source can be quiesced
before a destination file must be made available for read-write traffic.


|files_to_copy
|array[link:#files_to_copy[files_to_copy]]
a|A list of source files along with the destinations they are copied to.
If the terminal path component of the destination is a directory, then the
source file's basename is replicated in that directory.


|hold_quiescence
|boolean
a|Specifies whether the source file should be held quiescent for the
duration of the copy operation.


|max_throughput
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1MB/s is set to
1MB/s. A non-zero value greater than 1MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.
```

```
|reference_cutover_time
|integer
a|The maximum amount of time (in seconds) that the source reference file
can be quiesced before the corresponding destination file must be made
available for read-write traffic.


|reference_path
|string
a|The source reference file. If a reference file is specified, data for
other files being copied will be transferred as a difference from the
reference file. This can save bandwidth and destination storage if the
specified source files share blocks. If provided, this input must match
one of the source file paths. This input need not be provided if only one
source file is specified.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



= Manage storage file moves

:leveloffset: +1


[[IDcc2702972f450b5e0efe38407c313831]]
= Manage storage file moves



== Overview

You can use this API to start a file move operation between two FlexVol
volumes or within a FlexGroup volume, and view the status of all on-going
file move operations in the cluster.

&ndash; GET       /api/storage/file/moves

&ndash; GET       /api/storage/file/moves/{node.uuid}/\{uuid}/\{index}

&ndash; POST      /api/storage/file/moves

== Examples

=== Moving two files from one FlexVol volume to the other FlexVol volume

----

# The API:
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv.txt
move_between_fv.txt:
{
"files_to_move":{
  "sources":[
    {
```

```
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f1.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f2.txt"
    }
  ],
  "destinations":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/f1.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/f2.txt"
    }
  ]
}
}

# The response:
{ }
----

=== Moving two files from one FlexVol volume to the other FlexVol volume
```

```
(only specifying the destination directory)

----

# The API:
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv_dir.txt
move_between_fv_dir.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f3.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f4.txt"
    }
  ],
  "destinations":[
    {
      "svm":{
        "name":"vs0"
       },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/"
    }
  ]
}
}
```

```
# The response:
{ }
----

=== Moving multiple files from one FlexVol volume to the other FlexVol
volume and providing a source reference file

----

# The API:
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv_source.txt
move_between_fv_source.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f5.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f6.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f7.txt"
```

```
    }
  ],
  "destinations":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/f5.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/f6.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/f700.txt"
    }
  ]
},
"reference":{
  "svm":{
    "name":"vs0"
  },
  "volume":{
    "name":"fv1"
  },
  "path":"dir1/f6.txt"
}
}

# The response:
{ }
----
```

```
=== Moving a file between two FlexGroup volume constituents in the same
FlexGroup volume

----

# The API:
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fg.txt
move_between_fg.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fg2"
      },
      "path":"test/file.txt"
    }
  ],
  "destinations":[
    {
      "volume":{
        "name":"fg2__0008"
      }
    }
  ]
}
}

# The response:
{ }
----

=== Automatically selecting a destination constituent to move a file in a
FlexGroup volume for capacity rebalancing

----

# The API:
```

```
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/file/moves?automatic=true" -H
"accept: application/hal+json" -d @move_between_fg_automatic.txt
move_between_fg_automatic.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fg2"
      },
      "path":"test/file2.txt"
    }
  ]
}
}

# The response:
{ }
----

=== Moving a file between two FlexGroup volume constituents in the same
FlexGroup volume using the "force" and "disruptive" options

----

# The API:
/api/storage/file/moves

# The call:
curl -X POST  "https://<mgmt-
ip>/api/storage/file/moves?force=true&disruptive=true" -H "accept:
application/hal+json" -d @move_between_fg_force.txt
move_between_fg_force.txt
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
```

```
      "volume":{
        "name":"fg2"
      },
      "path":"test/file3.txt"
    }
  ],
  "destinations":[
    {
      "volume":{
        "name":"fg2__0008"
      }
    }
  ]
}
}

# The response:
{ }
----

=== Retrieving file move operations

----

# The API:
/api/storage/file/moves

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/file/moves" -H 'accept:
application/hal+json'

# The response:
{
"records": [
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "e12bc78d-36bb-4274-8163-fb8c21d59c9b",
```

```
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0008"
      },
      "path": "test/file2.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4d14f2f6-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0005"
      },
      "path": "test/file2.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/e12bc78d-36bb-4274-8163-fb8c21d59c9b/0"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
    "index": 1,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
```

```
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f2.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f2.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/1"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
```

```
        "name": "fv1"
      },
      "path": "dir1/f1.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f1.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/0"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "6d12601b-5377-43bf-99f0-b4bec37565e2",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f3.txt"
    },
```

```
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f3.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/6d12601b-5377-43bf-99f0-b4bec37565e2/0"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "6d12601b-5377-43bf-99f0-b4bec37565e2",
    "index": 1,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f4.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
```

```
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f4.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/6d12601b-5377-43bf-99f0-b4bec37565e2/1"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f5.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
```

```
      },
      "path": "dir2/f5.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/0"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
    "index": 2,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f7.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f700.txt"
    },
    "_links": {
```

```
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/2"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
    "index": 1,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f6.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f6.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/1"
      }
```

```
    }
  },
  {
    "node": {
      "uuid": "780255d2-10aa-11ec-a308-005056acf86d",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/780255d2-10aa-11ec-a308-
005056acf86d"
        }
      }
    },
    "uuid": "6591a42a-4ea2-4d40-bfb4-38959f6bd68e",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "438731dd-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0001"
      },
      "path": "test/file.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0008"
      },
      "path": "/test/file.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/780255d2-10aa-11ec-a308-
005056acf86d/6591a42a-4ea2-4d40-bfb4-38959f6bd68e/0"
      }
    }
  },
  {
    "node": {
```

```
      "uuid": "780255d2-10aa-11ec-a308-005056acf86d",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/780255d2-10aa-11ec-a308-
005056acf86d"
        }
      }
    },
    "uuid": "1a94e95a-346e-4eb3-969a-110e275cbf18",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "438731dd-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0001"
      },
      "path": "test/file3.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0008"
      },
      "path": "/test/file3.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/780255d2-10aa-11ec-a308-
005056acf86d/1a94e95a-346e-4eb3-969a-110e275cbf18/0"
      }
    }
  }
],
"num_records": 10,
"_links": {
  "self": {
    "href": "/api/storage/file/moves"
  }
```

```
}
}
----

=== Retrieving all moves in a file move operation

----

# The API:
/api/storage/file/moves/{node.uuid}/{uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/file/moves/*/ce2af347-586d-
4b31-b728-1e925f51fdfc/*" -H 'accept: application/hal+json'

# The response:
{
"records": [
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
    "index": 1,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f2.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
```

```
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
      },
      "path": "dir2/f2.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/1"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
        }
      }
    },
    "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv1"
      },
      "path": "dir1/f1.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
        "name": "fv2"
```

```
      },
      "path": "dir2/f1.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/0"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/file/moves/*/ce2af347-586d-4b31-b728-
1e925f51fdfc"
  }
}
}
----

=== Retrieving a specific file move in a file move operation

----

# The API:
/api/storage/file/moves/{node.uuid}/{uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/file/moves/*/ce2af347-586d-
4b31-b728-1e925f51fdfc/1" -H 'accept: application/hal+json'

# The response:
{
"node": {
  "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
  "name": "node1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
    }
  }
},
"uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
"index": 1,
"source": {
```

```
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv1"
  },
  "path": "dir1/f2.txt"
},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f2.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/1"
  }
}
}
----
```


```
[[ID229dd18f0371e4d27bc38d00155c6138]]
= Retrieve all ongoing file move operations in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/moves`#

*Introduced In:* 9.11

Retrieves all ongoing file move operations in the cluster.

== Related ONTAP commands

* `volume file move show`
```

```
* `volume rebalance file-move show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|is_flexgroup
|boolean
|query
|False
a|Filter by is_flexgroup


|is_destination_ready
|boolean
|query
|False
a|Filter by is_destination_ready


|is_snapshot_fenced
|boolean
|query
|False
a|Filter by is_snapshot_fenced
```

```
|failure.arguments.message
|string
|query
|False
a|Filter by failure.arguments.message


|failure.arguments.code
|string
|query
|False
a|Filter by failure.arguments.code


|failure.code
|string
|query
|False
a|Filter by failure.code


|failure.message
|string
|query
|False
a|Filter by failure.message


|index
|integer
|query
|False
a|Filter by index


|cutover_time
|integer
|query
|False
a|Filter by cutover_time


|scanner.percent
|integer
|query
```

```
|False
a|Filter by scanner.percent


|scanner.progress
|integer
|query
|False
a|Filter by scanner.progress


|scanner.state
|string
|query
|False
a|Filter by scanner.state


|scanner.total
|integer
|query
|False
a|Filter by scanner.total


|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|max_cutover_time
|integer
|query
|False
a|Filter by max_cutover_time


|source.path
```

```
|string
|query
|False
a|Filter by source.path


|source.svm.name
|string
|query
|False
a|Filter by source.svm.name


|source.svm.uuid
|string
|query
|False
a|Filter by source.svm.uuid


|source.volume.name
|string
|query
|False
a|Filter by source.volume.name


|source.volume.uuid
|string
|query
|False
a|Filter by source.volume.uuid


|elapsed_time
|integer
|query
|False
a|Filter by elapsed_time


|volume.name
|string
|query
|False
a|Filter by volume.name
```

```
|volume.uuid
|string
|query
|False
a|Filter by volume.uuid


|max_throughput
|integer
|query
|False
a|Filter by max_throughput


|uuid
|string
|query
|False
a|Filter by uuid


|destination.path
|string
|query
|False
a|Filter by destination.path


|destination.svm.name
|string
|query
|False
a|Filter by destination.svm.name


|destination.svm.uuid
|string
|query
|False
a|Filter by destination.svm.uuid


|destination.volume.name
|string
|query
|False
```

```
a|Filter by destination.volume.name


|destination.volume.uuid
|string
|query
|False
a|Filter by destination.volume.uuid


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records


|records
|array[link:#file_move[file_move]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "cutover_time": 8,
      "destination": {
        "path": "d1/d2/file1",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "elapsed_time": 100,
      "failure": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "index": 0,
      "max_cutover_time": 10,
      "max_throughput": 250000,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "scanner": {
      "percent": 80,
      "progress": 80000,
      "state": "string",
      "total": 100000
    },
    "source": {
      "path": "d1/d2/file1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
```

```
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
```

```
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
```

```
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#destination]
[.api-collapsible-fifth-title]
destination

Destination file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
```

```
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Contains the most recent failure reason for move operation.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#destinations]
[.api-collapsible-fifth-title]
destinations

Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|
```

```
|===


[#sources]
[.api-collapsible-fifth-title]
sources

Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#files_to_move]
[.api-collapsible-fifth-title]
files_to_move

A list of source files along with the destination file they are moved to.
If the terminal path component of the destination is a directory, then the
source file's basename is replicated in that directory. This is only used
for FlexVol volume file move operations.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|destinations
|array[link:#destinations[destinations]]
a|The destination file information.


|sources
|array[link:#sources[sources]]
a|The source file information.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node hosting the destination of this move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM of the source reference file.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

Volume of the source reference file.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
```

```
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#reference]
[.api-collapsible-fifth-title]
reference

Reference file information.

[#scanner]
[.api-collapsible-fifth-title]
scanner

[cols=3*,options=header]
|===
|Name
|Type
|Description

|percent
|integer
a|Scanner progress, as a percentage.


|progress
|integer
a|Scanner progress, in bytes scanned.


|state
|string
a|Status of the file move scanner.
```

```
|total
|integer
a|Total bytes to be scanned.


|===


[#source]
[.api-collapsible-fifth-title]
source

Source file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#file_move]
[.api-collapsible-fifth-title]
file_move

Details of a file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|Time that the file move operation takes before cutover completes, in
seconds.


|destination
|link:#destination[destination]
a|Destination file information.


|elapsed_time
|integer
a|Time elapsed since the start of the file move operation, in seconds.


|failure
|link:#error[error]
a|Contains the most recent failure reason for move operation.
```

```
|index
|integer
a|An additional unique element identifying one file among many that could
possibly be moved as part of a job. File index is an identifier ordered by
the file path arrays provided during create. For file move operations that
involve only one file, the file-index value of zero is always correct.


|is_destination_ready
|boolean
a|Indicates whether the destination file is ready for use.


|is_flexgroup
|boolean
a|Indicates whether this is a FlexGroup file move operation.


|is_snapshot_fenced
|boolean
a|Indicates whether snapshots are fenced.


|max_cutover_time
|integer
a|The maximum amount of time, in seconds that the source can be quiesced
before a destination file must be made available for read-write traffic.
Not supported in FlexGroup volume file move operations.


|max_throughput
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1 MB/s is set to 1
MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.


|node
|link:#node[node]
a|Node hosting the destination of this move operation.


|scanner
|link:#scanner[scanner]
```

```
a|

|source
|link:#source[source]
a|Source file information.


|svm
|link:#svm[svm]
a|SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


|uuid
|string
a|The UUID which uniquely identifies the job that started this move
operation.


|volume
|link:#volume[volume]
a|FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDa6cbfd774aec67dc2eddec1d4a8fe78d]]
= Start a file move operation

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/file/moves`#

*Introduced In:* 9.11

Starts a file move operation. This API can be used to move files from one
FlexVol volume to another FlexVol volume or within a FlexGroup volume for
capacity rebalancing.
For a FlexGroup volume file move operation, only one source file can be
specified in files_to_move. The source volume is the FlexGroup volume. The
destination volume is the destination FlexGroup volume constituent to move
the file to. When `automatic` is true, destination volume is not required.
The source path is the path to the file to be moved within the FlexGroup
volume. If the destination path is specified, it must be the same as the
source path.

== Required properties for file move operation

* `files_to_move` - List of files with the destination they are to be
moved to.

== Optional properties for file move operation

* `reference` - The source reference file for moving multiple files.
```

```
== Default property values

* `max_throughput` - _0_
* `max_cutover_time` - _10_
* `reference.max_cutover_time` - _10_

== Related ONTAP commands

* `volume file move start`
* `volume rebalance file-move start`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|force
|boolean
|query
|False
a|If true, the FlexGroup volume file move operation breaks the existing
lock state on the file being moved. Breaking the lock state may cause a
disruption for some client applications.

* Default value:


|automatic
|boolean
|query
|False
a|If true, the FlexGroup volume file move operation selects the
destination constituent automatically.

* Default value:


|disruptive
|boolean
|query
```

```
|False
a|If true, the FlexGroup volume file move operation is disruptive to the
clients. The file handle of the file being moved changes. If false, the
file handle remains the same. The non-disruptive file move operation is
only available on FlexGroup volumes with granular data property enabled.

* Default value:
* Introduced in: 9.12


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|Time that the file move operation takes before cutover completes, in
seconds.


|destination
|link:#destination[destination]
a|Destination file information.


|elapsed_time
|integer
a|Time elapsed since the start of the file move operation, in seconds.


|files_to_move
```

```
|link:#files_to_move[files_to_move]
a|A list of source files along with the destination file they are moved
to. If the terminal path component of the destination is a directory, then
the source file's basename is replicated in that directory. This is only
used for FlexVol volume file move operations.


|index
|integer
a|An additional unique element identifying one file among many that could
possibly be moved as part of a job. File index is an identifier ordered by
the file path arrays provided during create. For file move operations that
involve only one file, the file-index value of zero is always correct.


|is_destination_ready
|boolean
a|Indicates whether the destination file is ready for use.


|is_flexgroup
|boolean
a|Indicates whether this is a FlexGroup file move operation.


|is_snapshot_fenced
|boolean
a|Indicates whether snapshots are fenced.


|max_cutover_time
|integer
a|The maximum amount of time, in seconds that the source can be quiesced
before a destination file must be made available for read-write traffic.
Not supported in FlexGroup volume file move operations.


|max_throughput
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1 MB/s is set to 1
MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.


|node
```

```
|link:#node[node]
a|Node hosting the destination of this move operation.


|reference
|link:#reference[reference]
a|Reference file information.


|scanner
|link:#scanner[scanner]
a|

|source
|link:#source[source]
a|Source file information.


|svm
|link:#svm[svm]
a|SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


|uuid
|string
a|The UUID which uniquely identifies the job that started this move
operation.


|volume
|link:#volume[volume]
a|FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "cutover_time": 8,
  "destination": {
    "path": "d1/d2/file1",
```

```
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "elapsed_time": 100,
  "files_to_move": {
    "destinations": [
      {
        "path": "d1/d2/file1",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ],
    "sources": [
      {
        "path": "d1/d2/file1",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  },
  "index": 0,
  "max_cutover_time": 10,
  "max_throughput": 250000,
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "reference": {
```

```
    "max_cutover_time": 5,
    "path": "string",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "scanner": {
    "percent": 80,
    "progress": 80000,
    "state": "string",
    "total": 100000
  },
  "source": {
    "path": "d1/d2/file1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```


Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 918236
| The specified "volume.uuid" and "volume.name" do not refer to the same
volume.

| 2621462
| SVM `svm.name` does not exist.

| 2621706
| The specified "svm.uuid" and "svm.name" do not refer to the same SVM.

| 7012352
| File locations are inconsistent. All files must be on the same volume.

| 7012353
| Exceeded the file operations supported number of files.

| 7012354
| Unable to pair the number of source files to destination files.
```

```
| 7012357
| Cannot start a file operation until all cluster nodes support the file
operations capability.

| 7012358
| The specified source path is invalid.

| 7012359
| The specified destination path is invalid.

| 7012360
| The SVMs are not in an intracluster peering relationship.

| 7012361
| The SVMs peering relationship does not include application "file-move".

| 7012362
| The SVMs are not yet in a peered state.

| 7012363
| Cannot move files. All file operations must be managed by the
destination SVM's administrator.

| 7012364
| Cannot move file from SVM "svm.name" to SVM "svm.name". Moving a file
between SVMs is not supported.

| 7012365
| Copying a file between clusters is not supported.

| 7012367
| A reference path may only be specified if multiple source paths are
specified.

| 7012368
| The reference path must have a matching source path.

| 7012371
| The reference cutover time exceeds the maximum allowable time.

| 7012374
| Source volume and destination volume have different home clusters.

| 7012376
| Operation not allowed on a volume that is part of a SnapMirror
Synchronous relationship.
```

```
| 7012377
| Cannot start a file move operation on the volume because an active
volume conversion is in progress.

| 7013352
| One or more source files must be specified in "files_to_move.sources".

| 7013353
| One or more destination files or directories must be specified in
"files_to_move.destinations".

| 7013354
| "path" must be specified.

| 7013355
| Moving files between FlexVol volumes and FlexGroup volumes or
constituents is not supported.

| 7013356
| Moving files between FlexVol volumes and FlexGroup constituents is not
supported.

| 7013357
| The specified volume could not be found.

| 7013358
| The specified SVM or volume UUID could not be found.

| 7013359
| The SVM and volume must both be provided.

| 7018861
| Cannot start the file operation. The volume is offline.

| 7018877
| Maximum combined total (50) of file and LUN copy and move operations
reached. When one or more of the operations has completed, try the command
again.

| 7018937
| The file is already on the destination constituent.

| 13107222
| Internal error.

| 13107415
```

```
| Failed to lookup a volume property.

| 13107431
| Failed to lookup an SVM property.

| 13109260
| Failed to enable granular data on the volume.

| 144179201
| The file move start operation failed.

| 144179206
| Source file does not exist.

| 144179207
| Volume capacity balancing requires an effective cluster version of
9.10.1 or later.

| 144180200
| Destination constituent not a member of FlexGroup volume.

| 144180201
| Destination constituent not properly configured.

| 144180203
| Volume capacity rebalancing is not supported on FlexCache volumes.

| 144180204
| Volume capacity rebalancing is not supported on object store volumes.

| 144180205
| The system is busy.

| 144180207
| Volume capacity rebalancing is not supported on inactive MetroCluster
configurations.

| 144180208
| Disruptive file movement is not supported when granular data is enabled
on the volume. Try the operation again using "disruptive=false".

| 144181200
| Too many source or destination files are specified for a file move
within a FlexGroup volume. There must be one source file identifying a
file on a FlexGroup volume and either zero or one destination files
identifying the destination constituent.
```

```
| 144181201
| For a file move within a FlexGroup volume, the source volume must be a
FlexGroup volume, and the destination volume must be a constituent.

| 144181202
| The specified source volume UUID is for a constituent. For a file move
within a FlexGroup volume, the source volume must be a FlexGroup volume,
and the destination volume must be a constituent.

| 144181203
| A destination constituent must be provided in
"files_to_move.destinations" if it is not being selected automatically.
Use the "automatic" query to enable automatic destination constituent
selection.

| 144181204
| A destination constituent is provided while automatic destination
constituent selection is enabled with the "automatic" query.

| 144181205
| The destination volume is not a constituent. For a file move within a
FlexGroup volume, the destination volume must be a constituent of the
source FlexGroup volume.

| 144181207
| The destination constituent SVM is not the same as the source SVM. For a
file move within a FlexGroup volume, the destination constituent must be a
constituent of the source FlexGroup volume.

| 144181208
| The destination file path is different from the source file path. For a
file move within a FlexGroup volume, the path of the source file does not
change.

| 144181209
| The specified SVM for the destination constituent differs from the SVM
of the source FlexGroup volume. For a file move within a FlexGroup volume,
the destination constituent must be a constituent of the source FlexGroup
volume.

| 144181210
| The "force" parameter is not supported unless the "disruptive" parameter
is specified as "true".

| 144182201
| Volume capacity rebalancing using non-disruptive file move operations
```

```
and granular data requires an effective cluster version of 9.11.1 or
later.

| 144182228
| Operation is not supported on FlexGroup volumes with only one
constituent.

| 196608143
| Cannot start the operation. The volume is undergoing a secure purge
operation.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
* x-nullable: true


|===


[#destination]
[.api-collapsible-fifth-title]
destination

Destination file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
```

```
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Contains the most recent failure reason for move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#destinations]
[.api-collapsible-fifth-title]
destinations
```

```
Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#sources]
[.api-collapsible-fifth-title]
sources

Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.
```

```
|volume
|link:#volume[volume]
a|

|===


[#files_to_move]
[.api-collapsible-fifth-title]
files_to_move

A list of source files along with the destination file they are moved to.
If the terminal path component of the destination is a directory, then the
source file's basename is replicated in that directory. This is only used
for FlexVol volume file move operations.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destinations
|array[link:#destinations[destinations]]
a|The destination file information.


|sources
|array[link:#sources[sources]]
a|The source file information.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node hosting the destination of this move operation.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM of the source reference file.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

Volume of the source reference file.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#reference]
[.api-collapsible-fifth-title]
reference

Reference file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_cutover_time
|integer
a|The maximum amount of time, in seconds that the source reference file
can be quiesced before the corresponding destination file must be made
available for read-write traffic. Not supported in FlexGroup volume file
move operations.
```

```
|path
|string
a|The source reference file. If a reference file is specified, data for
other files being moved will be transferred as a difference from the
reference file. This can save bandwidth and destination storage if the
specified source files share blocks. If provided, this input must match
one of the source file paths. This input need not be provided if only one
source file is specified. Not supported in FlexGroup volume file move
operations.


|svm
|link:#svm[svm]
a|SVM of the source reference file.


|volume
|link:#volume[volume]
a|Volume of the source reference file.


|===


[#scanner]
[.api-collapsible-fifth-title]
scanner

[cols=3*,options=header]
|===
|Name
|Type
|Description

|percent
|integer
a|Scanner progress, as a percentage.


|progress
|integer
a|Scanner progress, in bytes scanned.


|state
|string
a|Status of the file move scanner.
```

```
|total
|integer
a|Total bytes to be scanned.


|===


[#source]
[.api-collapsible-fifth-title]
source

Source file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
* x-nullable: true


|===


[#file_move]
[.api-collapsible-fifth-title]
file_move

Details of a file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|Time that the file move operation takes before cutover completes, in
seconds.


|destination
|link:#destination[destination]
a|Destination file information.


|elapsed_time
|integer
a|Time elapsed since the start of the file move operation, in seconds.


|files_to_move
|link:#files_to_move[files_to_move]
a|A list of source files along with the destination file they are moved
to. If the terminal path component of the destination is a directory, then
the source file's basename is replicated in that directory. This is only
used for FlexVol volume file move operations.


|index
|integer
a|An additional unique element identifying one file among many that could
```

```
possibly be moved as part of a job. File index is an identifier ordered by
the file path arrays provided during create. For file move operations that
involve only one file, the file-index value of zero is always correct.


|is_destination_ready
|boolean
a|Indicates whether the destination file is ready for use.


|is_flexgroup
|boolean
a|Indicates whether this is a FlexGroup file move operation.


|is_snapshot_fenced
|boolean
a|Indicates whether snapshots are fenced.


|max_cutover_time
|integer
a|The maximum amount of time, in seconds that the source can be quiesced
before a destination file must be made available for read-write traffic.
Not supported in FlexGroup volume file move operations.


|max_throughput
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1 MB/s is set to 1
MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.


|node
|link:#node[node]
a|Node hosting the destination of this move operation.


|reference
|link:#reference[reference]
a|Reference file information.


|scanner
```

```
|link:#scanner[scanner]
a|

|source
|link:#source[source]
a|Source file information.


|svm
|link:#svm[svm]
a|SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


|uuid
|string
a|The UUID which uniquely identifies the job that started this move
operation.


|volume
|link:#volume[volume]
a|FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDbe6cf97090a18ba8dc542a2800ddeafc]]
= Check the status of an active file move operation

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/file/moves/{node.uuid}/{uuid}/{index}`#

*Introduced In:* 9.12

== Overview

Retrieve the status of an on-going file move operation.

== Related ONTAP commands

* `volume file move show`
* `volume rebalance file-move show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
```

```
|Description

|node.uuid
|string
|path
|True
a|Unique identifier of the node.


|uuid
|string
|path
|True
a|Unique identifier of the file move operation.


|index
|integer
|path
|True
a|The index of the file move within the overall operation


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_time
|integer
a|Time that the file move operation takes before cutover completes, in
seconds.
```

```
|destination
|link:#destination[destination]
a|Destination file information.


|elapsed_time
|integer
a|Time elapsed since the start of the file move operation, in seconds.


|failure
|link:#error[error]
a|Contains the most recent failure reason for move operation.


|index
|integer
a|An additional unique element identifying one file among many that could
possibly be moved as part of a job. File index is an identifier ordered by
the file path arrays provided during create. For file move operations that
involve only one file, the file-index value of zero is always correct.


|is_destination_ready
|boolean
a|Indicates whether the destination file is ready for use.


|is_flexgroup
|boolean
a|Indicates whether this is a FlexGroup file move operation.


|is_snapshot_fenced
|boolean
a|Indicates whether snapshots are fenced.


|max_cutover_time
|integer
a|The maximum amount of time, in seconds that the source can be quiesced
before a destination file must be made available for read-write traffic.
Not supported in FlexGroup volume file move operations.


|max_throughput
```

```
|integer
a|Maximum amount of data, in bytes that can be transferred per second in
support of this operation. A non-zero value less than 1 MB/s is set to 1
MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest
integral megabyte value. If unspecified, the default value is "0" which
means no range is set for the data transfer.


|node
|link:#node[node]
a|Node hosting the destination of this move operation.


|scanner
|link:#scanner[scanner]
a|

|source
|link:#source[source]
a|Source file information.


|svm
|link:#svm[svm]
a|SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


|uuid
|string
a|The UUID which uniquely identifies the job that started this move
operation.


|volume
|link:#volume[volume]
a|FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "cutover_time": 8,
  "destination": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "elapsed_time": 100,
  "failure": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "index": 0,
  "max_cutover_time": 10,
  "max_throughput": 250000,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
```

```
  "scanner": {
    "percent": 80,
    "progress": 80000,
    "state": "string",
    "total": 100000
  },
  "source": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
```

```
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 6521
| Cannot quiesce source file. Source file "file.name" is under
construction and is being used by either a file copy or move operation.
Wait a few minutes and try the operation again.

| 6522
| Volume is not VM aligned.

| 6523
| File suffix not VM aligned.

| 6524
| Completion interface not specified.

| 6525
| Maximum simultaneous file operations reached.

| 6526
| Lock conflict.

| 6527
| Permission denied.

| 6528
| Waiting for an I/O fence to be released.

| 6529
| Waiting for space to become available.

| 6530
| Waiting for an external process to complete.

| 6531
| Source volume was deleted.
```

```
| 6532
| Slowed due to resource constraints.

| 6533
| Volume does not support block reallocation.

| 6534
| File operation not supported.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
```

```
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#destination]
[.api-collapsible-fifth-title]
destination

Destination file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Contains the most recent failure reason for move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#destinations]
```

```
[.api-collapsible-fifth-title]
destinations

Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#sources]
[.api-collapsible-fifth-title]
sources

Path for file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|

|svm
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#files_to_move]
[.api-collapsible-fifth-title]
files_to_move

A list of source files along with the destination file they are moved to.
If the terminal path component of the destination is a directory, then the
source file's basename is replicated in that directory. This is only used
for FlexVol volume file move operations.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|destinations
|array[link:#destinations[destinations]]
a|The destination file information.


|sources
|array[link:#sources[sources]]
a|The source file information.


|===


[#node]
[.api-collapsible-fifth-title]
node

Node hosting the destination of this move operation.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM of the source reference file.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#volume]
[.api-collapsible-fifth-title]
volume

Volume of the source reference file.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#reference]
[.api-collapsible-fifth-title]
reference

Reference file information.
```

```
[#scanner]
[.api-collapsible-fifth-title]
scanner

[cols=3*,options=header]
|===
|Name
|Type
|Description

|percent
|integer
a|Scanner progress, as a percentage.


|progress
|integer
a|Scanner progress, in bytes scanned.


|state
|string
a|Status of the file move scanner.


|total
|integer
a|Total bytes to be scanned.


|===


[#source]
[.api-collapsible-fifth-title]
source

Source file information.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
```

```
|string
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM of the FlexGroup volume in file move operation. Only used in a
FlexGroup volume file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume
file move operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage FlexCache connection status

:leveloffset: +1


[[ID4b740a21de215737a731bbec418106af]]
= Manage FlexCache connection status



== Overview

FlexCache is a persistent cache of an origin volume.

FlexCache supports fan-out and more than one FlexCache can be created from
```

```
one origin volume.
This API retrieves the connection status between cache and origin volumes.

== FlexCache APIs

The following APIs can be used to perform operations related to FlexCache
connection status:

&ndash; GET       /api/storage/flexcache/connection-status

&ndash; GET       /api/storage/flexcache/connection-
status/\{node}/\{svm}/\{local_fg_msid}/\{remote_svm_uuid}/\{remote_vol_con
st_msid}

== Examples

=== Retrieving flexcache connection-status of all volumes.

The GET request is used to retrieve connection-status between all origins
and their corresponding FlexCache volumes.

----

# The API:
/api/storage/flexcache/connection-status

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/connection-status?"
-H  "accept: application/json"

# The response:
{
  "records": [
    {
      "node": "sti42-vsim-ucs511b",
      "svm": "vs0",
      "local_fg_msid": 2150359768,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359767,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511b/vs0/2150359768/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359767"
        }
      }
    },
    {
```

```
      "node": "sti42-vsim-ucs511d",
      "svm": "vs0",
      "local_fg_msid": 2150359768,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359767,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511d/vs0/2150359768/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359767"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359769,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359769"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359770,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359770"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359771,
      "_links": {
        "self": {
```

```
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359771"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359772,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359772"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359773,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359773"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359774,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359774"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
```

```
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359775,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359775"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359776,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359776"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359777,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359777"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359778,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
```

```
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359778"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359779,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359779"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359767,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359780,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359767/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359780"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511c",
      "svm": "vs0",
      "local_fg_msid": 2150359768,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359767,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359768/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359767"
        }
      }
    },
    {
      "node": "sti42-vsim-ucs511a",
      "svm": "vs0",
```

```
      "local_fg_msid": 2150359768,
      "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
      "remote_vol_const_msid": 2150359767,
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511a/vs0/2150359768/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359767"
        }
      }
    }
  ],
  "num_records": 16,
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/connection-status/"
    }
  }
}
----

=== Retrieving a single flexcache connection status

The GET request is used to retrieve a specific connection status between 2
volumes.

----

# The API:
/api/storage/flexcache/connection-
status/{node}/{svm}/{local_fg_msid}/{remote_svm_uuid}/{remote_vol_const_ms
id}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/connection-
status/sti42-vsim-ucs511c/vs0/2150359768/a9161fea-73a0-11f0-85c5-
005056a7bcdb/2150359767?" -H  "accept: application/json"

# The response:
{
  "node": "sti42-vsim-ucs511c",
  "svm": "vs0",
  "local_fg_msid": 2150359768,
  "remote_svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
  "remote_vol_const_msid": 2150359767,
  "svm_uuid": "a9161fea-73a0-11f0-85c5-005056a7bcdb",
  "volume": "fc0",
```

```
  "remote_svm": "vs0",
  "remote_volume": "vol0",
  "remote_cluster": "C1_sti42-vsim-ucs511a_1754576895",
  "conn_state": "connected",
  "remote_endpoint": "origin",
  "last_update_time": "2025-08-07T11:39:31-04:00",
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/connection-status/sti42-vsim-
ucs511c/vs0/2150359768/a9161fea-73a0-11f0-85c5-005056a7bcdb/2150359767"
    }
  }
}
```


```
----




[[ID643fd96feee77a1feac5181c436acb36]]
= Retrieve connection status between cache and origin volumes

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/connection-status`#

*Introduced In:* 9.18

Retrieves attributes of the connection status between cache and origin
volumes.

== Related ONTAP commands

* `volume flexcache connection-status show`

== Learn more

* xref:{relative_path}storage_flexcache_connection-
status_endpoint_overview.html[DOC /storage/flexcache/connection-status]
Retrieves origin of FlexCache in the cluster.


== Parameters

[cols=5*,options=header]
|===

|Name
```

```
|Type
|In
|Required
|Description

|local_fg_msid
|integer
|query
|False
a|Filter by local_fg_msid


|remote_vol_const_msid
|integer
|query
|False
a|Filter by remote_vol_const_msid


|node
|string
|query
|False
a|Filter by node


|remote_endpoint
|string
|query
|False
a|Filter by remote_endpoint


|remote_svm_uuid
|string
|query
|False
a|Filter by remote_svm_uuid


|last_update_time
|string
|query
|False
a|Filter by last_update_time
```

```
|svm
|string
|query
|False
a|Filter by svm


|remote_cluster
|string
|query
|False
a|Filter by remote_cluster


|svm_uuid
|string
|query
|False
a|Filter by svm_uuid


|conn_state
|string
|query
|False
a|Filter by conn_state


|volume
|string
|query
|False
a|Filter by volume


|remote_volume
|string
|query
|False
a|Filter by remote_volume


|remote_svm
|string
|query
|False
a|Filter by remote_svm
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records


|records
|array[link:#connection_status[connection_status]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
```

```
        }
      },
      "conn_state": "string",
      "last_update_time": "2018-06-04 19:00:00 +0000",
      "local_fg_msid": 0,
      "node": "string",
      "remote_cluster": "string",
      "remote_endpoint": "string",
      "remote_svm": "string",
      "remote_svm_uuid": "string",
      "remote_vol_const_msid": 0,
      "remote_volume": "string",
      "svm": "string",
      "svm_uuid": "string",
      "volume": "string"
    }
  ]
}
====

== Error
```


Status: Default, Error


```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
```

```
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#connection_status]
[.api-collapsible-fifth-title]
connection_status

FlexCache Connection Status Endpoint


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|conn_state
|string
a|Connection Status between the Flexcache and Origin volumes.
```

```
|last_update_time
|string
a|Last update time of the connection status.


|local_fg_msid
|integer
a|The local volume's Master Set ID.


|node
|string
a|Name of the node.


|remote_cluster
|string
a|Remote cluster.


|remote_endpoint
|string
a|Remote endpoint type.


|remote_svm
|string
a|Remote volume SVM.


|remote_svm_uuid
|string
a|Remote SVM UUID.


|remote_vol_const_msid
|integer
a|Remote volume's Master Set ID.


|remote_volume
|string
a|Remote Volume name.
```

```
|svm
|string
a|Volume SVM.


|svm_uuid
|string
a|Local SVM UUID.


|volume
|string
a|Local Volume name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID89f78f7ed3426b81615ddd21340f8b40]]
= Retrieve connection status between a cache and origin volume

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/connection-
status/{node}/{svm}/{local_fg_msid}/{remote_svm_uuid}/{remote_vol_const_ms
id}`#

*Introduced In:* 9.18

Retrieves connection status between a cache and origin volume.

== Related ONTAP commands

* `volume flexcache connection-status show`
```

```
== Learn more

* xref:{relative_path}storage_flexcache_connection-
status_endpoint_overview.html[DOC /storage/flexcache/connection-status]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node
|string
|path
|True
a|Node Name


|svm
|string
|path
|True
a|SVM Name


|local_fg_msid
|integer
|path
|True
a|Local FG MSID


|remote_svm_uuid
|string
|path
|True
a|Remote SVM UUID


|remote_vol_const_msid
```

```
|integer
|path
|True
a|Remote Constituent MSID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|conn_state
|string
a|Connection Status between the Flexcache and Origin volumes.


|last_update_time
|string
a|Last update time of the connection status.


|local_fg_msid
|integer
a|The local volume's Master Set ID.


|node
|string
a|Name of the node.
```

```
|remote_cluster
|string
a|Remote cluster.


|remote_endpoint
|string
a|Remote endpoint type.


|remote_svm
|string
a|Remote volume SVM.


|remote_svm_uuid
|string
a|Remote SVM UUID.


|remote_vol_const_msid
|integer
a|Remote volume's Master Set ID.


|remote_volume
|string
a|Remote Volume name.


|svm
|string
a|Volume SVM.


|svm_uuid
|string
a|Local SVM UUID.


|volume
|string
a|Local Volume name.
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "conn_state": "string",
  "last_update_time": "2018-06-04 19:00:00 +0000",
  "local_fg_msid": 0,
  "node": "string",
  "remote_cluster": "string",
  "remote_endpoint": "string",
  "remote_svm": "string",
  "remote_svm_uuid": "string",
  "remote_vol_const_msid": 0,
  "remote_volume": "string",
  "svm": "string",
  "svm_uuid": "string",
  "volume": "string"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#self_link]
[.api-collapsible-fifth-title]
self_link
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage FlexCache volumes

:leveloffset: +1


[[ID61f54191f4e407be896055617c3b95d7]]
= Manage FlexCache volumes



== Overview

FlexCache is a persistent cache of an origin volume. An origin volume can
only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:
```

```
&ndash; Intra-Vserver where FlexCache and the corresponding origin volume
reside in the same Vserver.

&ndash; Cross-Vserver but intra-cluster where FlexCache and the origin
volume reside in the same cluster but belong to different Vservers.

&ndash; Cross-cluster where FlexCache and the origin volume reside in
different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from
one origin volume.
This API retrieves and manages FlexCache configurations in the cache
cluster.

== FlexCache APIs

The following APIs can be used to perform operations related with
FlexCache:

&ndash; GET       /api/storage/flexcache/flexcaches

&ndash; GET       /api/storage/flexcache/flexcaches/\{uuid}

&ndash; POST      /api/storage/flexcache/flexcaches

&ndash; DELETE    /api/storage/flexcache/flexcaches/\{uuid}

== Examples

=== Creating a FlexCache

The POST request is used to create a FlexCache.

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H  "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\",
\"origins\": [ {  \"svm\": { \"name\": \"vs_3\"  }, \"volume\": {
\"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" } }"

# The response:
{
```

```
"job": {
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H  "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\",
\"origins\": [ {  \"svm\": { \"name\": \"vs_3\"  }, \"volume\": {
\"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" },  \"path\":
\"/fc_333\", \"prepopulate\": { \"dir_paths\": [ \"/dir1\" ] } }"

# The response:
{
"job": {
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H  "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\":           \"aggr_1\" } ], \"name\":
```

```
\"fc_333\", \"origins\": [ {  \"svm\": { \"name\": \"vs_3\"  },
\"volume\": { \"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" },
\"path\": \"/       fc_333\", \"prepopulate\": { \"dir_paths\": [
\"/dir1\" ], \"exclude_dir_paths\": [ \"/dir1/dir11\" ] } }"

# The response:
{
"job": {
  "uuid": "5afe9ea4-1dcf-11eb-b006-005056ac6a93",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/5afe9ea4-1dcf-11eb-b006-005056ac6a93"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H  "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\",
\"origins\": [ {  \"svm\": { \"name\": \"vs_3\"  }, \"volume\": {
\"name\": \"vol_o1\" } } ], \"svm\":{ \"name\": \"vs_1\" }, \"dr_cache\":
true,  \"path\": \"/fc_333\", \"prepopulate\": { \"dir_paths\": [
\"/dir1\" ] } }"

# The response:
{
"job": {
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----

=== Retrieving FlexCache attributes
```

```
The GET request is used to retrieve FlexCache attributes. The object
includes a large set of fields which can be expensive to retrieve. Most
notably, the fields size, guarantee.type, aggregates, path,
origins.ip_address, origins.size, and origins.state are expensive to
retrieve. The recommended method to use this API is to filter and retrieve
only the required fields.

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/flexcaches?" -H
"accept: application/json"

# The response:
{
"records": [
  {
    "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79",
    "name": "fc_322",
    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches/04d5e07b-0ebe-11e9-
8180-0050568e0b79"
      }
    }
  },
  {
    "uuid": "47902654-0ea4-11e9-8180-0050568e0b79",
    "name": "fc_321",
    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches/47902654-0ea4-11e9-
8180-0050568e0b79"
      }
    }
  },
  {
    "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79",
    "name": "fc_323",
    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches/77e911ff-0ebe-11e9-
8180-0050568e0b79"
```

```
      }
    }
  },
  {
    "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79",
    "name": "fc_32",
    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches/ddb42bbc-0e95-11e9-
8180-0050568e0b79"
      }
    }
  },
  {
    "uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79",
    "name": "fc_333",
    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-
8b2b-0050568e0b79"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/flexcache/flexcaches?"
  }
}
}
----

=== Retrieving the attributes of a FlexCache

The GET request is used to retrieve the attributes of a FlexCache. The
object includes a large set of fields which can be expensive to retrieve.
Most notably, the fields size, guarantee.type, aggregates, path,
origins.ip_address, origins.size, and origins.state are expensive to
retrieve. The recommended method to use this API is to filter and retrieve
only the required fields.

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/flexcaches/ec774932-
0f3c-11e9-8b2b-0050568e0b79" -H  "accept: application/json"

# The response:
{
"uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79",
"name": "fc_333",
"svm": {
  "name": "vs_1",
  "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
},
"size": 4294967296,
"guarantee": {
    "type": "volume"
},
"dr_cache": "true",
"aggregates": [
  {
    "name": "aggr_1",
    "uuid": "26f34b76-88f8-4a47-b5e0-d8e901fb1114"
  }
],
"origins": [
  {
    "ip_address": "10.140.103.175",
    "size": 20971520,
    "create_time": "2019-01-03T15:19:55+05:30",
    "state": "online",
    "volume": {
      "name": "vol_o1",
      "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313"
    },
    "svm": {
      "name": "vs_3",
      "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
    },
    "cluster": {
      "name": "node2",
      "uuid": "50733f81-0e90-11e9-b391-0050568e4115"
    }
  }
],
"_links": {
  "self": {
```

```
    "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-
0050568e0b79"
  }
}
}
----

=== Deleting a FlexCache

The DELETE request is used to delete a FlexCache.

----

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H  "accept: application/json"

# The response:
{
"job": {
  "uuid": "e17994f2-0f3e-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e17994f2-0f3e-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----

=== Modifying a FlexCache volume

Use the PATCH request to update a FlexCache volume.

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H  "accept: application/json" -H  "Content-Type: application/json" -d "{
```

```
\"prepopulate\": { \"dir_paths\": [ \"/dir1\" ] } }"

# The response:
{
"job": {
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H  "accept: application/json" -H  "Content-Type: application/json" -d "{
\"prepopulate\": { \"dir_paths\": [ \"/dir1\" ], \"exclude_dir_paths\": [
\"/dir1/dir11\" ] } }"

# The response:
{
"job": {
  "uuid": "b574c48c-1da7-11eb-b006-005056ac6a93",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b574c48c-1da7-11eb-b006-005056ac6a93"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
```

```
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H  "Content-Type: application/json" -d "{
\"nfsv4\" : { \"enabled\": true } }"

# The response:
{
"job": {
  "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"nfsv4\" : { \"enabled\": false } }"

# The response:
{
"job": {
  "uuid": "17e193f3-304b-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/17e193f3-304b-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}
```

```
# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H  "Content-Type: application/json" -d "{
\"cifs\" : { \"enabled\": true } }"

# The response:
{
"job": {
  "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"cifs\" : { \"enabled\": false } }"

# The response:
{
"job": {
  "uuid": "17e193f3-304b-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/17e193f3-304b-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
```

```
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H  "Content-Type: application/json" -d "{
\"s3\" : { \"enabled\": true } }"

# The response:
{
"job": {
  "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H  "Content-Type: application/json" -d "{
\"s3\" : { \"enabled\": false } }"

# The response:
{
"job": {
  "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----
```

```
# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H  "Content-Type: application/json" -d "{
\"writeback\" : { \"enabled\": true } }"

# The response:
{
"job": {
  "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"writeback\" : { \"enabled\": false } }"

# The response:
{
"job": {
  "uuid": "17e193f3-304b-11ed-a5d5-005056bbb2b7",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/17e193f3-304b-11ed-a5d5-005056bbb2b7"
    }
  }
}
}
----
```

```
----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"relative_size" : {"enabled":true,
"percentage" : 50}}'

# The response:
{
"job": {
  "uuid": "34d9c90d-ac38-11ed-a2ca-005056bb5573",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/34d9c90d-ac38-11ed-a2ca-005056bb5573"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"relative_size" : {"enabled":false}}'

# The response:
{
"job": {
  "uuid": "14bfcc28-ac82-11ed-83bd-005056bb5573",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/14bfcc28-ac82-11ed-83bd-005056bb5573"
    }
  }
}
}
----
```

```
----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"atime_scrub" : {"enabled":true,
"period" : 30}}'

# The response:
{
"job": {
  "uuid": "34d9c90d-ac38-11ed-a2ca-005056bb5573",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/34d9c90d-ac38-11ed-a2ca-005056bb5573"
    }
  }
}
}
----

----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"atime_scrub" : {"enabled":false}}'

# The response:
{
"job": {
  "uuid": "14bfcc28-ac82-11ed-83bd-005056bb5573",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/14bfcc28-ac82-11ed-83bd-005056bb5573"
    }
  }
}
}
----
```

```
----

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"cifs_change_notify" :
{"enabled":true}}'

# The response:
{
"job": {
  "uuid": "14bfcc28-ac82-11ed-83bd-005056bb5573",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/14bfcc28-ac82-11ed-83bd-005056bb5573"
    }
  }
}
}
----




[[ID8a465ea6fe62533719649b8f4ce91fde]]
= Retrieve a FlexCache volume in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/flexcaches`#

*Introduced In:* 9.6

Retrieves FlexCache in the cluster.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `origins.ip_address` - IP address of origin.
```

```
* `origins.size` - Physical size of origin.
* `origins.state` - State of origin.
* `size` - Physical size of FlexCache.
* `guarantee.type` - Space guarantee style of FlexCache.
* `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of
FlexCache volume.
* `path` - Fully-qualified path of the owning SVM's namespace where the
FlexCache is mounted.

== Related ONTAP commands

* `volume flexcache show`

== Learn more

*
xref:{relative_path}storage_flexcache_flexcaches_endpoint_overview.html[DO
C /storage/flexcache/flexcaches]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|relative_size.enabled
|boolean
|query
|False
a|Filter by relative_size.enabled

* Introduced in: 9.13


|relative_size.percentage
|integer
|query
|False
a|Filter by relative_size.percentage

* Introduced in: 9.13
```

```
|origins.create_time
|string
|query
|False
a|Filter by origins.create_time


|origins.ip_address
|string
|query
|False
a|Filter by origins.ip_address


|origins.state
|string
|query
|False
a|Filter by origins.state


|origins.volume.name
|string
|query
|False
a|Filter by origins.volume.name


|origins.volume.uuid
|string
|query
|False
a|Filter by origins.volume.uuid


|origins.svm.name
|string
|query
|False
a|Filter by origins.svm.name


|origins.svm.uuid
|string
|query
```

```
|False
a|Filter by origins.svm.uuid


|origins.size
|integer
|query
|False
a|Filter by origins.size


|origins.cluster.name
|string
|query
|False
a|Filter by origins.cluster.name


|origins.cluster.uuid
|string
|query
|False
a|Filter by origins.cluster.uuid


|s3.enabled
|boolean
|query
|False
a|Filter by s3.enabled

* Introduced in: 9.18


|atime_scrub.period
|integer
|query
|False
a|Filter by atime_scrub.period

* Introduced in: 9.14


|atime_scrub.enabled
|boolean
|query
|False
```

```
a|Filter by atime_scrub.enabled

* Introduced in: 9.14


|global_file_locking_enabled
|boolean
|query
|False
a|Filter by global_file_locking_enabled

* Introduced in: 9.9


|use_tiered_aggregate
|boolean
|query
|False
a|Filter by use_tiered_aggregate

* Introduced in: 9.8


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|override_encryption
|boolean
|query
|False
a|Filter by override_encryption

* Introduced in: 9.14


|constituents_per_aggregate
```

```
|integer
|query
|False
a|Filter by constituents_per_aggregate


|guarantee.type
|string
|query
|False
a|Filter by guarantee.type

* Introduced in: 9.7


|dr_cache
|boolean
|query
|False
a|Filter by dr_cache

* Introduced in: 9.9


|uuid
|string
|query
|False
a|Filter by uuid


|size
|integer
|query
|False
a|Filter by size


|nfsv4.enabled
|boolean
|query
|False
a|Filter by nfsv4.enabled

* Introduced in: 9.18
```

```
|cifs_change_notify.enabled
|boolean
|query
|False
a|Filter by cifs_change_notify.enabled

* Introduced in: 9.14


|aggregates.name
|string
|query
|False
a|Filter by aggregates.name


|aggregates.uuid
|string
|query
|False
a|Filter by aggregates.uuid


|name
|string
|query
|False
a|Filter by name

* maxLength: 203
* minLength: 1


|writeback.enabled
|boolean
|query
|False
a|Filter by writeback.enabled

* Introduced in: 9.12


|cifs.enabled
|boolean
|query
|False
a|Filter by cifs.enabled
```

```
* Introduced in: 9.18


|path
|string
|query
|False
a|Filter by path


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#flexcache[flexcache]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "constituents_per_aggregate": 1,
      "guarantee": {
        "type": "string"
      },
      "name": "vol1",
      "origins": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "create_time": "2018-06-04 19:00:00 +0000",
          "ip_address": "10.10.10.7",
          "size": 0,
          "state": "string",
          "svm": {
            "_links": {
              "self": {
```

```
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "path": "/user/my_fc",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#atime_scrub]
[.api-collapsible-fifth-title]
atime_scrub

FlexCache volume atime-based scrub


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether scrubbing of inactive files based on atime is enabled
for the FlexCache volume. When scrubbing is enabled, files whose atime is
older than the specified duration are evicted from the cache volume. The
scrubber runs once a day and looks for files whose atime has exceeded the
provisioned value.


|period
|integer
a|Specifies the atime duration in days after which a cached file is
considered inactive. Inactive files are purged from the FlexCache volumes
when the scrubber runs once a day.
```

```
|===


[#cifs]
[.api-collapsible-fifth-title]
cifs

FlexCache CIFS


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether CIFS access is enabled on the FlexCache volume.


|===


[#cifs_change_notify]
[.api-collapsible-fifth-title]
cifs_change_notify

Notifies that a change has been made to the FlexCache volume CIFS.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether a CIFS change notification is enabled for the
FlexCache volume.


|===
```

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee

[cols=3*,options=header]
|===
|Name
|Type
|Description

|type
|string
a|The type of space guarantee of this volume in the aggregate.


|===


[#nfsv4]
[.api-collapsible-fifth-title]
nfsv4

FlexCache NFSv4


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether NFSv4 access is enabled on the FlexCache volume.


|===


[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
```

```
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#prepopulate]
[.api-collapsible-fifth-title]
prepopulate

FlexCache prepopulate

[#relative_size]
[.api-collapsible-fifth-title]
relative_size

FlexCache Relative Size
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether the relative sizing is enabled for the FlexCache
volume. Relative sizing is introduced as a part of follow the origin
feature. When relative sizing is enabled, it blocks any modifications done
manually in the absolute size of the FlexCache volume. The size of the
FlexCache volume is calculated and entered automatically based on the size
of the origin volume.


|percentage
|integer
a|Specifies the percent size the FlexCache volume should have relative to
the total size of the origin volume. This property is only relevant to a
FlexCache volume that has the relative size property enabled.


|===


[#s3]
[.api-collapsible-fifth-title]
s3

Flexcache S3


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether S3 access is enabled on the Flexcache volume.


|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

FlexCache SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#writeback]
[.api-collapsible-fifth-title]
writeback

FlexCache Writeback


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
```

```
a|Indicates whether or not writeback is enabled for the FlexCache volume.
Writeback is a storage method where data is first written to the FlexCache
volume and then written to the origin of a FlexCache volume.


|===


[#flexcache]
[.api-collapsible-fifth-title]
flexcache

Defines the cache endpoint of FlexCache.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|aggregates
|array[link:#aggregates[aggregates]]
a|

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub


|cifs
|link:#cifs[cifs]
a|FlexCache CIFS


|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|constituents_per_aggregate
|integer
a|Number of FlexCache constituents per aggregate when the 'aggregates'
```

```
field is mentioned.


|dr_cache
|boolean
a|If set to true, a DR cache is created.


|global_file_locking_enabled
|boolean
a|Specifies whether or not a FlexCache volume has global file locking mode
enabled. Global file locking mode is a mode where protocol read locking
semantics are enforced across all FlexCaches and origins of a FlexCache
volume. When global file locking mode is enabled, the
"is_disconnected_mode_off_for_locks" flag is always set to "true".


|guarantee
|link:#guarantee[guarantee]
a|

|name
|string
a|FlexCache name


|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4


|origins
|array[link:#flexcache_relationship[flexcache_relationship]]
a|

|override_encryption
|boolean
a|If set to true, a plaintext FlexCache volume for an encrypted origin
volume is created.


|path
|string
a|The fully-qualified path in the owning SVM's namespace at which the
FlexCache is mounted. The path is case insensitive and must be unique
within a SVM's namespace. Path must begin with '/' and must not end with
'/'. Only one FlexCache be mounted at any given junction path.
```

```
|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|size
|integer
a|Physical size of the FlexCache. The recommended size for a FlexCache is
10% of the origin volume. The minimum FlexCache constituent size is 1GB.


|svm
|link:#svm[svm]
a|FlexCache SVM


|use_tiered_aggregate
|boolean
a|Specifies whether or not a Fabricpool-enabled aggregate can be used in
FlexCache creation. The use_tiered_aggregate is only used when auto-
provisioning a FlexCache volume.


|uuid
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID3b1f7c17cc2956d7372434c0dd841f2a]]
= Create a FlexCache volume in the cluster

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/flexcache/flexcaches`#

*Introduced In:* 9.6

Creates a FlexCache in the cluster.

== Required properties

* `name` - Name of FlexCache volume.
* `origins.volume.name` - Name of the origin volume. This volume can only
be identified by its name, not by its UUID.
* `origins.svm.name` - Name of origin Vserver.
* `svm.name` or `svm.uuid` - Name or UUID of Vserver where FlexCache will
be created.

== Recommended optional properties

* `path` - Path to mount the FlexCache volume
* `prepopulate.dir_paths` - List of directory-paths to be prepopulated for
the FlexCache volume.
* `prepopulate.exclude_dir_paths` - List of directory-paths to be excluded
from prepopulation for the FlexCache volume.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `size` - 10% of origin volume size or 1GB per constituent, whichever is
greater.
* `guarantee.type` - none. FlexCache is thin provisioned by default.
* `constituents_per_aggregate` - 4 if aggregates.name or aggregates.uuid
is used.
```

```
* `use_tiered_aggregate` - false if aggr-list is not used. This property
is only used when auto-provisioning a FlexCache volume.
* `is_disconnected_mode_off_for_locks` - false. This property specifies if
the origin will honor the cache side locks when doing the lock checks in
the disconnected mode.
* `dr_cache` - false if FlexCache is not a DR cache. This property is used
to create a DR FlexCache.
* `global_file_locking_enabled` - false. This property specifies whether
global file locking is enabled on the FlexCache volume.
* `writeback.enabled` - false. This property specifies whether writeback
is enabled for the FlexCache volume.
* `relative_size.enabled` - false. This property specifies whether the
relative sizing is enabled for the FlexCache volume.
* `relative_size.percentage` - 10. This property specifies the percent
size FlexCache volume should have relative to the total size of the origin
volume.
* `override_encryption` - false. If true, this property is used to create
a plaintext FlexCache volume for an encrypted origin volume.
* `atime_scrub.enabled` - false. This property specifies whether scrubbing
of inactive files based on atime is enabled for the FlexCache volume.
* `atime_scrub.period` - 30. This property specifies the atime duration in
days after which the file can be scrubbed from the FlexCache volume if it
stays unused beyond the duration.
* `nfsv4.enabled` - false. This property specifies whether NFSv4 is
enabled for the FlexCache volume.
* `cifs.enabled` - false. This property specifies whether CIFS is enabled
for the FlexCache volume.
* `s3.enabled` - false. This property specifies whether S3 is enabled for
the FlexCache volume.
* `cifs_change_notify.enabled` - false. This property specifies whether a
CIFS change notification is enabled for the FlexCache volume.

== Related ONTAP commands

* `volume flexcache create`
* `volume flexcache prepopulate start`

== Learn more

*
xref:{relative_path}storage_flexcache_flexcaches_endpoint_overview.html[DO
C /storage/flexcache/flexcaches]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
```

```
|array[link:#aggregates[aggregates]]
a|

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub


|cifs
|link:#cifs[cifs]
a|FlexCache CIFS


|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|constituents_per_aggregate
|integer
a|Number of FlexCache constituents per aggregate when the 'aggregates'
field is mentioned.


|dr_cache
|boolean
a|If set to true, a DR cache is created.


|global_file_locking_enabled
|boolean
a|Specifies whether or not a FlexCache volume has global file locking mode
enabled. Global file locking mode is a mode where protocol read locking
semantics are enforced across all FlexCaches and origins of a FlexCache
volume. When global file locking mode is enabled, the
"is_disconnected_mode_off_for_locks" flag is always set to "true".


|guarantee
|link:#guarantee[guarantee]
a|

|name
|string
a|FlexCache name
```

```
|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4


|origins
|array[link:#flexcache_relationship[flexcache_relationship]]
a|

|override_encryption
|boolean
a|If set to true, a plaintext FlexCache volume for an encrypted origin
volume is created.


|path
|string
a|The fully-qualified path in the owning SVM's namespace at which the
FlexCache is mounted. The path is case insensitive and must be unique
within a SVM's namespace. Path must begin with '/' and must not end with
'/'. Only one FlexCache be mounted at any given junction path.


|prepopulate
|link:#prepopulate[prepopulate]
a|FlexCache prepopulate


|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|size
|integer
a|Physical size of the FlexCache. The recommended size for a FlexCache is
10% of the origin volume. The minimum FlexCache constituent size is 1GB.


|svm
|link:#svm[svm]
a|FlexCache SVM
```

```
|use_tiered_aggregate
|boolean
a|Specifies whether or not a Fabricpool-enabled aggregate can be used in
FlexCache creation. The use_tiered_aggregate is only used when auto-
provisioning a FlexCache volume.


|uuid
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "constituents_per_aggregate": 1,
  "guarantee": {
    "type": "string"
  },
  "name": "vol1",
  "origins": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "ip_address": "10.10.10.7",
```

```
      "size": 0,
      "state": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "path": "/user/my_fc",
  "prepopulate": {
    "dir_paths": [
      "/dir1"
    ],
    "exclude_dir_paths": [
      "/dir1/dir2"
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11
| Provided value is smaller than the required minimum flexvolume size.

| 262186
| Invalid combination with field "aggregates".

| 917937
| The root path "/" is an invalid value for the field "path".

| 918723
| Aggregate value cannot be used to create a volume on this platform.

| 66846735
| The peer state of the relationship is not peered or the FlexCache
application is not supported by the relationship

| 66846762
| The origin volume in the SVM is offline

| 66846767
| Volume does not exist in SVM.

| 66846867
| Global file locking (GFL) mode is already disabled for origin of a
FlexCache volume in SVM. Enabling GFL is only supported when the first
FlexCache volume is created for an origin. Set
is_global_file_locking_enabled" to "false".

| 66846768
| Volume already exists in SVM.

| 66846863
| The "preserve-msid" parameter is not supported when creating the
FlexCache volume in the same SVM as the origin of a FlexCache volume

| 66846870
```

```
| Either the SVM name or origin volume name is missing

| 66846871
| Constituents per aggregate are specified but aggregate name is missing

| 66846872
| More than one origin volume is specified

| 66846873
| The specified SVM UUID is incorrect for the specified SVM name

| 66846874
| The specified name and SVM uuid refer to different aggregates.

| 66846875
| The specified aggregate name does not exist

| 66846876
| The specified SVM does not exist or is not peered

| 66846877
| The specified origin SVM name is of zero length

| 66846878
| The specified SVM UUID is invalid

| 66846998
| Failed to modify the "relative_size_percentage" property of volume in
SVM because the "relative_size_enabled" property is false.

| 66846730
| Failed to create a FlexCache volume

| 66846760
| The specified SVM is not a data Vserver

| 66846787
| The specified aggregate is a SnapLock aggregate

| 66846812
| The specified aggregate is a Composite aggregate

| 66846812
| The specified junction path is under a FlexCache volume

| 66846834
| FlexCache encryption requires a cluster version of 9.6 or higher
```

```
| 66846835
| A volume encryption license is not found

| 66847013
| Creating a FlexCache volume with override-encryption property set to
true requires an effective cluster version of 9.14.0 or later on both
origin and cache clusters

| 66846844
| An object store server volume cannot be the origin of a FlexCache volume

| 66846902
| FlexCache is not supported on an All SAN Array system.

| 66846915
| The "use_tiered_aggregate" option is only supported when auto
provisioning the FlexCache volume.

| 66847026
| Failed to enable the atime-scrub-enabled property for the FlexCache
volume because the atime-update-property is false
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#atime_scrub]
[.api-collapsible-fifth-title]
atime_scrub

FlexCache volume atime-based scrub


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether scrubbing of inactive files based on atime is enabled
for the FlexCache volume. When scrubbing is enabled, files whose atime is
older than the specified duration are evicted from the cache volume. The
```

```
scrubber runs once a day and looks for files whose atime has exceeded the
provisioned value.


|period
|integer
a|Specifies the atime duration in days after which a cached file is
considered inactive. Inactive files are purged from the FlexCache volumes
when the scrubber runs once a day.


|===


[#cifs]
[.api-collapsible-fifth-title]
cifs

FlexCache CIFS


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether CIFS access is enabled on the FlexCache volume.


|===


[#cifs_change_notify]
[.api-collapsible-fifth-title]
cifs_change_notify

Notifies that a change has been made to the FlexCache volume CIFS.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Specifies whether a CIFS change notification is enabled for the
FlexCache volume.


|===


[#guarantee]
[.api-collapsible-fifth-title]
guarantee

[cols=3*,options=header]
|===
|Name
|Type
|Description

|type
|string
a|The type of space guarantee of this volume in the aggregate.


|===


[#nfsv4]
[.api-collapsible-fifth-title]
nfsv4

FlexCache NFSv4


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether NFSv4 access is enabled on the FlexCache volume.


|===
```

```
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|
```

```
|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#prepopulate]
[.api-collapsible-fifth-title]
prepopulate

FlexCache prepopulate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|dir_paths
|array[string]
```

```
a|

|exclude_dir_paths
|array[string]
a|

|recurse
|boolean
a|Specifies whether or not the prepopulate action should search through
the `dir_paths` recursively. If not set, the default value _true_ is used.


|===


[#relative_size]
[.api-collapsible-fifth-title]
relative_size

FlexCache Relative Size


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether the relative sizing is enabled for the FlexCache
volume. Relative sizing is introduced as a part of follow the origin
feature. When relative sizing is enabled, it blocks any modifications done
manually in the absolute size of the FlexCache volume. The size of the
FlexCache volume is calculated and entered automatically based on the size
of the origin volume.


|percentage
|integer
a|Specifies the percent size the FlexCache volume should have relative to
the total size of the origin volume. This property is only relevant to a
FlexCache volume that has the relative size property enabled.


|===
```

```
[#s3]
[.api-collapsible-fifth-title]
s3

Flexcache S3


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether S3 access is enabled on the Flexcache volume.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

FlexCache SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===
```

```
[#writeback]
[.api-collapsible-fifth-title]
writeback

FlexCache Writeback


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether or not writeback is enabled for the FlexCache volume.
Writeback is a storage method where data is first written to the FlexCache
volume and then written to the origin of a FlexCache volume.


|===


[#flexcache]
[.api-collapsible-fifth-title]
flexcache

Defines the cache endpoint of FlexCache.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub
```

```
|cifs
|link:#cifs[cifs]
a|FlexCache CIFS


|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|constituents_per_aggregate
|integer
a|Number of FlexCache constituents per aggregate when the 'aggregates'
field is mentioned.


|dr_cache
|boolean
a|If set to true, a DR cache is created.


|global_file_locking_enabled
|boolean
a|Specifies whether or not a FlexCache volume has global file locking mode
enabled. Global file locking mode is a mode where protocol read locking
semantics are enforced across all FlexCaches and origins of a FlexCache
volume. When global file locking mode is enabled, the
"is_disconnected_mode_off_for_locks" flag is always set to "true".


|guarantee
|link:#guarantee[guarantee]
a|

|name
|string
a|FlexCache name


|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4


|origins
|array[link:#flexcache_relationship[flexcache_relationship]]
a|
```

```
|override_encryption
|boolean
a|If set to true, a plaintext FlexCache volume for an encrypted origin
volume is created.


|path
|string
a|The fully-qualified path in the owning SVM's namespace at which the
FlexCache is mounted. The path is case insensitive and must be unique
within a SVM's namespace. Path must begin with '/' and must not end with
'/'. Only one FlexCache be mounted at any given junction path.


|prepopulate
|link:#prepopulate[prepopulate]
a|FlexCache prepopulate


|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|size
|integer
a|Physical size of the FlexCache. The recommended size for a FlexCache is
10% of the origin volume. The minimum FlexCache constituent size is 1GB.


|svm
|link:#svm[svm]
a|FlexCache SVM


|use_tiered_aggregate
|boolean
a|Specifies whether or not a Fabricpool-enabled aggregate can be used in
FlexCache creation. The use_tiered_aggregate is only used when auto-
provisioning a FlexCache volume.
```

```
|uuid
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID527fadd686fcc30447925f1587aa1da8]]
= Delete a FlexCache volume

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/flexcache/flexcaches/{uuid}`#

*Introduced In:* 9.6

Deletes a FlexCache. If a FlexCache volume is online, it is offlined
before deletion.

== Related ONTAP commands

* `volume flexcache delete`

== Learn more

*
xref:{relative_path}storage_flexcache_flexcaches_endpoint_overview.html[DO
C /storage/flexcache/flexcaches]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Unique identifier of the FlexCache.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
```

```
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66846879
| The specified volume UUID is not a FlexCache volume

| 66846731
| Failed to delete the FlexCache volume

| 524546
| Failed to delete the FlexCache volume because the FlexCache volume is
not unmounted

| 66846980
| Failed to delete the FlexCache volume because it has the writeback
property enabled
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[ID34c5f423fb79873001bce1d03daf8ad3]]
= Retrieve attributes of the FlexCache volume in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/flexcaches/{uuid}`#

*Introduced In:* 9.6

Retrieves attributes of the FlexCache in the cluster.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are included by default in GET. The recommended method to
use this API is to filter and retrieve only the required fields. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `origins.ip_address` - IP address of origin.
* `origins.size` - Physical size of origin.
* `origins.state` - State of origin.
* `size` - Physical size of FlexCache.
* `guarantee.type` - Space guarantee style of FlexCache.
* `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of
FlexCache volume.
* `path` - Fully-qualified path of the owning SVM's namespace where the
FlexCache is mounted.

== Related ONTAP commands

* `volume flexcache show`

== Learn more

*
xref:{relative_path}storage_flexcache_flexcaches_endpoint_overview.html[DO
C /storage/flexcache/flexcaches]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|uuid
|string
|path
|True
a|Unique identifier of FlexCache.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|aggregates
|array[link:#aggregates[aggregates]]
a|

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub


|cifs
|link:#cifs[cifs]
a|FlexCache CIFS
```

```
|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|constituents_per_aggregate
|integer
a|Number of FlexCache constituents per aggregate when the 'aggregates'
field is mentioned.


|dr_cache
|boolean
a|If set to true, a DR cache is created.


|global_file_locking_enabled
|boolean
a|Specifies whether or not a FlexCache volume has global file locking mode
enabled. Global file locking mode is a mode where protocol read locking
semantics are enforced across all FlexCaches and origins of a FlexCache
volume. When global file locking mode is enabled, the
"is_disconnected_mode_off_for_locks" flag is always set to "true".


|guarantee
|link:#guarantee[guarantee]
a|

|name
|string
a|FlexCache name


|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4


|origins
|array[link:#flexcache_relationship[flexcache_relationship]]
a|

|override_encryption
|boolean
a|If set to true, a plaintext FlexCache volume for an encrypted origin
volume is created.
```

```
|path
|string
a|The fully-qualified path in the owning SVM's namespace at which the
FlexCache is mounted. The path is case insensitive and must be unique
within a SVM's namespace. Path must begin with '/' and must not end with
'/'. Only one FlexCache be mounted at any given junction path.


|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|size
|integer
a|Physical size of the FlexCache. The recommended size for a FlexCache is
10% of the origin volume. The minimum FlexCache constituent size is 1GB.


|svm
|link:#svm[svm]
a|FlexCache SVM


|use_tiered_aggregate
|boolean
a|Specifies whether or not a Fabricpool-enabled aggregate can be used in
FlexCache creation. The use_tiered_aggregate is only used when auto-
provisioning a FlexCache volume.


|uuid
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "constituents_per_aggregate": 1,
  "guarantee": {
    "type": "string"
  },
  "name": "vol1",
  "origins": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#atime_scrub]
```

```
[.api-collapsible-fifth-title]
atime_scrub

FlexCache volume atime-based scrub


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether scrubbing of inactive files based on atime is enabled
for the FlexCache volume. When scrubbing is enabled, files whose atime is
older than the specified duration are evicted from the cache volume. The
scrubber runs once a day and looks for files whose atime has exceeded the
provisioned value.


|period
|integer
a|Specifies the atime duration in days after which a cached file is
considered inactive. Inactive files are purged from the FlexCache volumes
when the scrubber runs once a day.


|===


[#cifs]
[.api-collapsible-fifth-title]
cifs

FlexCache CIFS


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether CIFS access is enabled on the FlexCache volume.
```

```
|===


[#cifs_change_notify]
[.api-collapsible-fifth-title]
cifs_change_notify

Notifies that a change has been made to the FlexCache volume CIFS.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether a CIFS change notification is enabled for the
FlexCache volume.


|===


[#guarantee]
[.api-collapsible-fifth-title]
guarantee

[cols=3*,options=header]
|===
|Name
|Type
|Description

|type
|string
a|The type of space guarantee of this volume in the aggregate.


|===


[#nfsv4]
[.api-collapsible-fifth-title]
```

```
nfsv4

FlexCache NFSv4


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether NFSv4 access is enabled on the FlexCache volume.


|===


[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#prepopulate]
[.api-collapsible-fifth-title]
prepopulate

FlexCache prepopulate

[#relative_size]
[.api-collapsible-fifth-title]
relative_size

FlexCache Relative Size


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether the relative sizing is enabled for the FlexCache
volume. Relative sizing is introduced as a part of follow the origin
feature. When relative sizing is enabled, it blocks any modifications done
manually in the absolute size of the FlexCache volume. The size of the
FlexCache volume is calculated and entered automatically based on the size
of the origin volume.


|percentage
|integer
a|Specifies the percent size the FlexCache volume should have relative to
the total size of the origin volume. This property is only relevant to a
```

```
FlexCache volume that has the relative size property enabled.


|===


[#s3]
[.api-collapsible-fifth-title]
s3

Flexcache S3


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether S3 access is enabled on the Flexcache volume.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

FlexCache SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#writeback]
[.api-collapsible-fifth-title]
writeback

FlexCache Writeback


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether or not writeback is enabled for the FlexCache volume.
Writeback is a storage method where data is first written to the FlexCache
volume and then written to the origin of a FlexCache volume.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====
```

```
[[IDf788ffff3163795089823803d71c7636]]
= Pre-populate a FlexCache volume in the cluster

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/flexcache/flexcaches/{uuid}`#

*Introduced In:* 9.8

Prepopulates a FlexCache volume in the cluster, or modifies configuration
of the FlexCache volume.

== Recommended optional properties

* `uuid` - FlexCache volume UUID.
* `prepopulate.exclude_dir_paths` - List of directory-paths to be excluded
from prepopulation for the FlexCache volume.
* `prepopulate.dir_paths` - List of directory-paths to be prepopulated for
the FlexCache volume.
* `writeback.enabled` - false. This property specifies whether writeback
is enabled for the FlexCache volume.
* `relative_size.enabled` - This property specifies whether the relative
sizing is enabled for the FlexCache volume.
* `relative_size.percentage` - This property specifies the percent size
FlexCache volume should have relative to the total size of the origin
volume.
* `atime_scrub.enabled` - This property specifies whether the atime based
scrub is enabled for the FlexCache volume.
* `atime_scrub.period` - This property specifies the duration in days
after which inactive files can be scrubbed from FlexCache volume.
* `cifs_change_notify.enabled` - This property specifies whether a CIFS
change notification is enabled for the FlexCache volume.
* `nfsv4.enabled` - This property specifies whether NFSv4 is enabled for
the FlexCache volume.
* `cifs.enabled` - This property specifies whether CIFS is enabled for the
FlexCache volume.
* `s3.enabled` - This property specifies whether S3 is enabled for the
FlexCache volume.

== Default property values

If not specified in PATCH, the following default property value is
assigned:

* `prepopulate.recurse` - Default value is "true".
```

```
== Related ONTAP commands

* `volume flexcache prepopulate start`
* `volume flexcache config modify`

== Learn more

*
xref:{relative_path}storage_flexcache_flexcaches_endpoint_overview.html[DO
C /storage/flexcache/flexcaches]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Unique identifier of the FlexCache volume.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub


|cifs
|link:#cifs[cifs]
a|FlexCache CIFS


|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4


|prepopulate
|link:#prepopulate[prepopulate]
a|FlexCache prepopulate


|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|uuid
```

```
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "dr_cache": true,
  "global_file_locking_enabled": true,
  "override_encryption": true,
  "prepopulate": {
    "dir_paths": [
      "/dir1"
    ],
    "exclude_dir_paths": [
      "/dir1/dir2"
    ]
  },
  "use_tiered_aggregate": true,
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66846921
| Failed to prepopulate. Empty path list

| 66846922
| FlexCache volume does not exist in the SVM

| 66846923
```

```
| Failed to prepopulate FlexCache volume because the origin volume is not
reachable

| 66846924
| FlexCache volume is offline

| 66846925
| FlexCache volume is not mounted or the junction-path is not active

| 66846926
| The junction-path of FlexCache volume is not active

| 66846927
| FlexCache volume does not have an origin volume

| 66846928
| FlexCache volume does not exist

| 66846929
| Using FlexCache prepopulate requires an effective cluster version of
9.8.0 or later

| 66846930
| Using FlexCache prepopulate in a MetroCluster configuration requires an
effective cluster version of 9.8.0 or later on both the local and remote
clusters

| 66846931
| Internal Error. FlexCache prepopulate job queue failed. Wait a few
minutes, and then try the operation again

| 66846936
| Failed to lookup root file handle for origin of FlexCache volume. Wait a
few minutes, and then try the operation again

| 66846937
| Internal error. Failed to initialize thread

| 66846939
| Internal error. Failed to get the MSID of the origin volume for
FlexCache volume

| 66846943
| Failed to prepopulate because dir_path does not exist

| 66846944
| Failed to get root constituent for FlexCache volume
```

```
| 66846945
| Origin of FlexCache volume is not mounted or the junction-path is not
active

| 66846946
| The junction-path of origin of FlexCache volume is not active

| 66846947
| FlexCache prepopulate job for FlexCache volume already exists

| 66846948
| FlexCache prepopulate job for FlexCache volume could not be queued
because the node is offline

| 66846958
| Exclude path(s) does not exist in origin volume

| 66846988
| Cannot enable writeback for FlexCache volume because the origin is a
data protection volume. Writeback is only supported on FlexCache volumes
with origins that are read-write volumes

| 66846995
| Using FlexCache relative-size requires an effective cluster version of
9.13.1 or later

| 66846997
| Failed to modify the relative-size-percentage property of the FlexCache
volume because the used size of the FlexCache Volume is greater than the
newly calculated size based on the value of the relative-size-percentage
property

| 66846997
| Failed to modify the relative-size-percentage property of FlexCache
volume because the relative-size-percentage property is not valid

| 66846998
| Failed to modify the relative-size-percentage property of FlexCache
volume because the is-relative-size-enabled property is false

| 66847010
| Failed to modify the relative-size-percentage property of FlexCache
volume because autosize is enabled on the FlexCache volume

| 66847026
| Failed to enable the atime-scrub-enabled property for the FlexCache
```

```
volume because the atime-update-property is false

| 66847027
| Failed to modify atime-scrub-period property of the FlexCache volume
because the atime-scrub-enabled property is false

| 66846720
| Using FlexCache nfsv4 access requires an effective cluster version of
9.10.0 or later

| 66846720
| Using FlexCache cifs access requires an effective cluster version of
9.8.0 or later

| 66847093
| Using FlexCache S3 access requires an effective cluster version of
9.18.1 or later
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#atime_scrub]
[.api-collapsible-fifth-title]
atime_scrub

FlexCache volume atime-based scrub


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether scrubbing of inactive files based on atime is enabled
for the FlexCache volume. When scrubbing is enabled, files whose atime is
older than the specified duration are evicted from the cache volume. The
scrubber runs once a day and looks for files whose atime has exceeded the
provisioned value.


|period
```

```
|integer
a|Specifies the atime duration in days after which a cached file is
considered inactive. Inactive files are purged from the FlexCache volumes
when the scrubber runs once a day.


|===


[#cifs]
[.api-collapsible-fifth-title]
cifs

FlexCache CIFS


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether CIFS access is enabled on the FlexCache volume.


|===


[#cifs_change_notify]
[.api-collapsible-fifth-title]
cifs_change_notify

Notifies that a change has been made to the FlexCache volume CIFS.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether a CIFS change notification is enabled for the
FlexCache volume.
```

```
|===


[#guarantee]
[.api-collapsible-fifth-title]
guarantee

[cols=3*,options=header]
|===
|Name
|Type
|Description

|type
|string
a|The type of space guarantee of this volume in the aggregate.


|===


[#nfsv4]
[.api-collapsible-fifth-title]
nfsv4

FlexCache NFSv4


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether NFSv4 access is enabled on the FlexCache volume.


|===


[#cluster]
[.api-collapsible-fifth-title]
cluster
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
```

```
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.
```

```
|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state


|volume
|link:#volume[volume]
a|

|===


[#prepopulate]
[.api-collapsible-fifth-title]
prepopulate

FlexCache prepopulate


[cols=3*,options=header]
|===
|Name
|Type
|Description

|dir_paths
|array[string]
a|

|exclude_dir_paths
|array[string]
a|

|recurse
|boolean
a|Specifies whether or not the prepopulate action should search through
the `dir_paths` recursively. If not set, the default value _true_ is used.
```

```
|===


[#relative_size]
[.api-collapsible-fifth-title]
relative_size

FlexCache Relative Size


[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether the relative sizing is enabled for the FlexCache
volume. Relative sizing is introduced as a part of follow the origin
feature. When relative sizing is enabled, it blocks any modifications done
manually in the absolute size of the FlexCache volume. The size of the
FlexCache volume is calculated and entered automatically based on the size
of the origin volume.


|percentage
|integer
a|Specifies the percent size the FlexCache volume should have relative to
the total size of the origin volume. This property is only relevant to a
FlexCache volume that has the relative size property enabled.


|===


[#s3]
[.api-collapsible-fifth-title]
s3

Flexcache S3


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether S3 access is enabled on the Flexcache volume.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

FlexCache SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#writeback]
[.api-collapsible-fifth-title]
writeback

FlexCache Writeback


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether or not writeback is enabled for the FlexCache volume.
Writeback is a storage method where data is first written to the FlexCache
volume and then written to the origin of a FlexCache volume.


|===


[#flexcache]
[.api-collapsible-fifth-title]
flexcache

Defines the cache endpoint of FlexCache.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|atime_scrub
|link:#atime_scrub[atime_scrub]
a|FlexCache volume atime-based scrub


|cifs
|link:#cifs[cifs]
a|FlexCache CIFS


|cifs_change_notify
|link:#cifs_change_notify[cifs_change_notify]
a|Notifies that a change has been made to the FlexCache volume CIFS.


|nfsv4
|link:#nfsv4[nfsv4]
a|FlexCache NFSv4
```

```
|prepopulate
|link:#prepopulate[prepopulate]
a|FlexCache prepopulate


|relative_size
|link:#relative_size[relative_size]
a|FlexCache Relative Size


|s3
|link:#s3[s3]
a|Flexcache S3


|uuid
|string
a|FlexCache UUID. Unique identifier for the FlexCache.


|writeback
|link:#writeback[writeback]
a|FlexCache Writeback


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage FlexCache origin volumes

:leveloffset: +1


[[ID238a26d4a6ba4fa962408321df488843]]
= Manage storage FlexCache origin volumes



== Overview

FlexCache is a persistent cache of an origin volume. An origin volume can
only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:

&ndash; Intra-Vserver where FlexCache and the corresponding origin volume
reside in the same Vserver.

&ndash; Cross-Vserver but intra-cluster where FlexCache and the origin
volume reside in the same cluster but belong to different Vservers.

&ndash; Cross-cluster where FlexCache and the origin volume reside in
different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from
one origin volume.
This API retrieves the origin of FlexCache configurations in the origin
```

```
cluster.

== FlexCache APIs

The following APIs can be used to perform operations related to the origin
of a FlexCache:

&ndash; GET       /api/storage/flexcache/origins

&ndash; GET       /api/storage/flexcache/origins/\{uuid}

&ndash; PATCH     /api/storage/flexcache/origins/\{uuid}

== Examples

=== Retrieving origins of FlexCache attributes

The GET request is used to retrieve the origins of FlexCache attributes.

----

# The API:
/api/storage/flexcache/origins

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins?" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313",
      "name": "vol_o1",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/2bc957dd-2617-4afb-8d2f-
66ac6070d313"
        }
      }
    },
    {
      "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
      "name": "vol_1",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-
```

```
86d28a34cda8"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/origins?"
    }
  }
}
----

=== Retrieving the attributes of an origin volume

The GET request is used to retrieve the attributes of an origin volume.

----

# The API:
/api/storage/flexcache/origins/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins/80fcaee4-
0dc2-488b-afb8-86d28a34cda8" -H  "accept: application/json"

# The response:
{
  "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
  "name": "vol_1",
  "svm": {
    "name": "vs_3",
    "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
  },
  "block_level_invalidation": "false",
  "global_file_locking_enabled": "true",
  "flexcaches": [
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T19:27:22+05:30",
      "volume": {
        "name": "fc_42",
        "uuid": "4e7f9d49-0e96-11e9-aed0-0050568eddbe"
      },
      "svm": {
```

```
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T21:08:34+05:30",
      "volume": {
        "name": "fc_421",
        "uuid": "71ee8f36-0ea4-11e9-aed0-0050568eddbe"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-03T11:14:38+05:30",
      "volume": {
        "name": "fc_422"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-02T19:24:14+05:30",
      "state": "online",
      "volume": {
        "name": "fc_32",
```

```
        "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-02T21:07:23+05:30",
      "state": "online",
      "volume": {
        "name": "fc_321",
        "uuid": "47902654-0ea4-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-03T00:11:38+05:30",
      "state": "online",
      "volume": {
        "name": "fc_322",
        "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
```

```
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-03T00:14:52+05:30",
      "state": "online",
      "volume": {
        "name": "fc_323",
        "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-
86d28a34cda8"
    }
  }
}
----

=== Modifying origin options of an origin volume

Use the PATCH request to update options of an origin volume.

----

# The API:
/api/storage/flexcache/origins/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/flexcache/origins/1fbc0ebb-
2440-11eb-a86c-005056ac8ca0"  -H  "accept: application/json" -H  "Content-
Type: application/json" -d "{ \"block_level_invalidation\": \"true\" } "

# The response:
{
"job": {
```

```
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
----
```


```
[[ID3ec371139194e0aea1e02b15b6457de4]]
= Retrieve the origin of a FlexCache in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/origins`#

*Introduced In:* 9.6

Retrieves origin of FlexCache in the cluster.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `flexcaches.ip_address` - IP address of FlexCache.
* `flexcaches.size` - Physical size of FlexCache.
* `flexcaches.guarantee.type` - Space guarantee style of FlexCache.
* `flexcaches.state` - State of FlexCache.

== Related ONTAP commands

* `volume flexcache origin show-caches`

== Learn more

* xref:{relative_path}storage_flexcache_origins_endpoint_overview.html[DOC
/storage/flexcache/origins]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|query
|False
a|Filter by uuid


|block_level_invalidation
|boolean
|query
|False
a|Filter by block_level_invalidation

* Introduced in: 9.9


|global_file_locking_enabled
|boolean
|query
|False
a|Filter by global_file_locking_enabled

* Introduced in: 9.9


|name
|string
|query
|False
a|Filter by name

* maxLength: 203
* minLength: 1


|flexcaches.create_time
```

```
|string
|query
|False
a|Filter by flexcaches.create_time


|flexcaches.ip_address
|string
|query
|False
a|Filter by flexcaches.ip_address


|flexcaches.state
|string
|query
|False
a|Filter by flexcaches.state


|flexcaches.volume.name
|string
|query
|False
a|Filter by flexcaches.volume.name


|flexcaches.volume.uuid
|string
|query
|False
a|Filter by flexcaches.volume.uuid


|flexcaches.svm.name
|string
|query
|False
a|Filter by flexcaches.svm.name


|flexcaches.svm.uuid
|string
|query
|False
a|Filter by flexcaches.svm.uuid
```

```
|flexcaches.size
|integer
|query
|False
a|Filter by flexcaches.size


|flexcaches.cluster.name
|string
|query
|False
a|Filter by flexcaches.cluster.name


|flexcaches.cluster.uuid
|string
|query
|False
a|Filter by flexcaches.cluster.uuid


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
```

```
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#flexcache_origin[flexcache_origin]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "flexcaches": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
```

```
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "create_time": "2018-06-04 19:00:00 +0000",
          "ip_address": "10.10.10.7",
          "size": 0,
          "state": "string",
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "name": "vol1, vol_2",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
Origin volume SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#flexcache_origin]
[.api-collapsible-fifth-title]
flexcache_origin

Defines the origin endpoint of FlexCache.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|block_level_invalidation
|boolean
a|Block level invalidation enables the FlexCache volume to retain blocks
```

```
that are not changed at the FlexCache volume without having to evict them.
This means that the FlexCache volume does not have to again incur the
computational cost of fetching blocks over the WAN from the FlexCache
volume origin on the next client access. Block level invalidation is a
property of the origin volume. Without block level invalidation, any write
at the origin volume would evict the whole file at the FlexCache volume,
since by default, origin volume does a file level invalidation.


|flexcaches
|array[link:#flexcache_relationship[flexcache_relationship]]
a|

|global_file_locking_enabled
|boolean
a|Specifies whether a global file locking option is enabled for an origin
volume of a FlexCache volume.


|name
|string
a|Origin volume name


|svm
|link:#svm[svm]
a|Origin volume SVM


|uuid
|string
a|Origin volume UUID. Unique identifier for origin of FlexCache.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
```

```
====



[[IDa1741a5dfdf23f5418d6468120e5e6ce]]
= Retrieve attributes of the origin of a FlexCache volume in the cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/flexcache/origins/{uuid}`#

*Introduced In:* 9.6

Retrieves attributes of the origin of a FlexCache in the cluster.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are included by default in GET results. The recommended
method to use this API is to filter and retrieve only the required fields.
See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `flexcaches.ip_address` - IP address of FlexCache.
* `flexcaches.size` - Physical size of FlexCache.
* `flexcaches.guarantee.type` - Space guarantee style of FlexCache.
* `flexcaches.state` - State of FlexCache.
* `flexcaches.dr_cache` - True if the cache is a DR cache.

== Related ONTAP commands

* `volume flexcache origin show-caches`

== Learn more

* xref:{relative_path}storage_flexcache_origins_endpoint_overview.html[DOC
/storage/flexcache/origins]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
```

```
|Required
|Description

|uuid
|string
|path
|True
a|Unique identifier of origin of FlexCache.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|block_level_invalidation
|boolean
a|Block level invalidation enables the FlexCache volume to retain blocks
that are not changed at the FlexCache volume without having to evict them.
This means that the FlexCache volume does not have to again incur the
computational cost of fetching blocks over the WAN from the FlexCache
volume origin on the next client access. Block level invalidation is a
property of the origin volume. Without block level invalidation, any write
at the origin volume would evict the whole file at the FlexCache volume,
since by default, origin volume does a file level invalidation.


|flexcaches
|array[link:#flexcache_relationship[flexcache_relationship]]
```

```
a|

|global_file_locking_enabled
|boolean
a|Specifies whether a global file locking option is enabled for an origin
volume of a FlexCache volume.


|name
|string
a|Origin volume name


|svm
|link:#svm[svm]
a|Origin volume SVM


|uuid
|string
a|Origin volume UUID. Unique identifier for origin of FlexCache.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "flexcaches": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
```

```
      "create_time": "2018-06-04 19:00:00 +0000",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "name": "vol1, vol_2",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.
```

```
|state
|string
a|Volume state


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

Origin volume SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID7aeb380b9bc8cd0a633554ae6a89f6bf]]
= Modify origin options for a Flexcache origin volume in the cluster

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/flexcache/origins/{uuid}`#

*Introduced In:* 9.9

Modifies origin options for a origin volume in the cluster.

== Required properties

* `uuid` - Origin volume UUID.
* `block_level_invalidation` - Value for the Block Level Invalidation flag
- options {true\|false}.

== Related ONTAP commands

* `volume flexcache origin config modify`

== Learn more

* xref:{relative_path}storage_flexcache_origins_endpoint_overview.html[DOC
/storage/flexcache/origins]


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Unique identifier of the origin of the FlexCache volume.


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_level_invalidation
|boolean
a|Block level invalidation enables the FlexCache volume to retain blocks
that are not changed at the FlexCache volume without having to evict them.
This means that the FlexCache volume does not have to again incur the
```

```
computational cost of fetching blocks over the WAN from the FlexCache
volume origin on the next client access. Block level invalidation is a
property of the origin volume. Without block level invalidation, any write
at the origin volume would evict the whole file at the FlexCache volume,
since by default, origin volume does a file level invalidation.


|global_file_locking_enabled
|boolean
a|Specifies whether a global file locking option is enabled for an origin
volume of a FlexCache volume.


|uuid
|string
a|Origin volume UUID. Unique identifier for origin of FlexCache.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 66846870
| Either origin volume uuid or BLI option value is missing

| 66847020
| Failed to get origin volume details using volume uuid

| 66847021
```

```
| Failed to modify origin volume options
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
```

```
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#flexcache_relationship]
[.api-collapsible-fifth-title]
flexcache_relationship

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|create_time
|string
a|Creation time of the relationship.


|ip_address
|string
a|Cluster management IP of the remote cluster.


|size
|integer
a|Size of the remote volume.


|state
|string
a|Volume state
```

```
|volume
|link:#volume[volume]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

Origin volume SVM


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#flexcache_origin]
[.api-collapsible-fifth-title]
flexcache_origin

Defines the origin endpoint of FlexCache.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|block_level_invalidation
|boolean
a|Block level invalidation enables the FlexCache volume to retain blocks
that are not changed at the FlexCache volume without having to evict them.
This means that the FlexCache volume does not have to again incur the
computational cost of fetching blocks over the WAN from the FlexCache
volume origin on the next client access. Block level invalidation is a
property of the origin volume. Without block level invalidation, any write
at the origin volume would evict the whole file at the FlexCache volume,
since by default, origin volume does a file level invalidation.


|global_file_locking_enabled
|boolean
a|Specifies whether a global file locking option is enabled for an origin
volume of a FlexCache volume.


|uuid
|string
a|Origin volume UUID. Unique identifier for origin of FlexCache.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage pools

:leveloffset: +1


[[ID0e170fa2fecb7738e7d29528ac80b55f]]
= Manage storage pools



== Retrieving storage pool information

The Storage Pools GET API retrieves all shared storage pools in the
cluster.

The collection GET returns the storage pool identifiers, UUID and name.
The instance GET, by default, returns all of the properties defined in the
storage_pool object.
The "show_spares" query returns a response outside of the records body,
which includes the groups of usable spares in the cluster.
The usable count for each class of spares does not include reserved spare
capacity recommended by ONTAP best practices.

== Creating storage pools

Creating a shared storage pool is recommended when distributing flash
capacity across the cache tiers of HDD aggregates across an HA pair.
POST can be used with specific properties to create a storage pool as
requested. At a minimum, the
storage pool name, disk count, and the nodes where it should reside, are
```

```
required to create a new instance.

When using POST with input properties, three properties are required.
These are:

* name - Name of the storage pool.
* node.name or node.uuid - Node that can use capacity from the storage
pool in their cache tiers.
* capacity.disk_count - Number of disks to be used to create the storage
pool.

== Examples

=== Retrieving a list of storage pools from the cluster

The following example shows the response with a list of storage pools in
the cluster:

----
{
"records": [
{
  "uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
  "nodes": [
    {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
    {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/pools/8255fef7-4737-11ec-bd1b-005056bbb879"
    }
  }
}
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/pools"
  }
}
```

```
}
----

----

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools?fields=*" -H "accept:
application/json"

# The response:
{
"records": [
{
  "uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
  "name": "new_sp",
  "nodes": [
    {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
    {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    }
  ],
  "storage_type": "ssd",
  "capacity": {
      "remaining": 1846542336,
      "total": 7386169344,
      "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "count": 1,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 1846542336
      },
      {
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
```

```
          "name": "node-2",
        },
        "count": 0,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 0
      }
      ],
      "used_allocation_units": [
      {
        "aggregate": {
          "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
          "name": "test_a"
        },
        "allocated_unit_count": 2,
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "capacity": 2769813504
      },
      {
        "aggregate": {
          "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
          "name": "test_b"
        },
        "allocated_unit_count": 1,
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
          "name": "node-2",
        },
        "capacity": 1384906752
      }
      ],
      "disk_count": 4,
      "disks": [
        {
          "disk": {
            "name": "VMw-1.11"
          },
          "usable_size": 1902379008,
          "total_size": 1908871168,
        },
        {
          "disk": {
            "name": "VMw-1.12"
```

```
          },
          "usable_size": 1902379008,
          "total_size": 1908871168,
        },
        {
          "disk": {
            "name": "VMw-1.23"
          },
          "usable_size": 1902379008,
          "total_size": 1908871168,
        },
        {
          "disk": {
            "name": "VMw-1.24"
          },
          "usable_size": 1902379008,
          "total_size": 1908871168,
        }
      ]
    },
    "health": {
      "state": "normal",
      "is_healthy": true
    },
}
],
"num_records": 1,
}
```


---

=== Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare information according to ONTAP best practices.

---

```
 # The API:
/api/storage/pools?show_spares=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools?show_spares=true" -H
"accept: application/json"

# The response:
{
```

```
"records": [],
"num_records": 0,
"spares": [
  {
    "disk_class": "solid_state",
    "disk_type": "ssd",
    "size": 3720609792,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "usable": 12,
    "nodes": [
      {
        "uuid": "54af4069-c1f7-11ec-884e-005056bb6e0b",
        "name": "node-1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/54af4069-c1f7-11ec-884e-
005056bb6e0b"
          }
        }
      },
      {
        "uuid": "d50f1acb-c1f6-11ec-9dfd-005056bb8d04",
        "name": "node-2",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/d50f1acb-c1f6-11ec-9dfd-
005056bb8d04"
          }
        }
      }
    ]
  }
],
"_link": {
  "self": {
    "href": "/api/storage/aggregates?show_spares=true"
  }
}
}
----

=== Simulating the creation of a storage pool

The following example shows the response containing the simulated layout
details of a new storage pool in the cluster.
```

```
NOTE: Each storage pool UUID provided in this response is not guaranteed
to be the same UUID for the storage pool if it is created.

----

# The API:
/api/storage/pools

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/pools?simulate=true&fields=*"
-d "{\"nodes\": [{\"name\": \"node1\"}, {\"name\": \"node2\"}], \"name\":
\"storage_pool_1\", \"capacity\": {\"disk_count\": \"4\"}}" -H "accept:
application/json"

# The response:
{
"records": [
{
  "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
  "name": "new_sp",
  "nodes": [
    {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
    {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    }
  ],
  "capacity": {
    "total": 7386169344,
    "disk_count": 4,
    "disks": [
      {
        "disk": {
          "name": "VMw-1.11"
        },
      },
      {
        "disk": {
          "name": "VMw-1.12"
        },
      },
      {
```

```
        "disk": {
          "name": "VMw-1.23"
        },
      },
      {
        "disk": {
          "name": "VMw-1.24"
        },
      }
    ]
  }
}
]
}
```


---

```
[[ID39767ed2076e9dc56c77a4abc33d73d1]]
= Retrieve storage pools for the entire cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/pools`#

*Introduced In:* 9.11

Retrieves the collection of storage pools for the entire cluster.

== Related ONTAP commands

* `storage pool show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
```

```
|string
|query
|False
a|Filter by name


|capacity.disks.disk.name
|string
|query
|False
a|Filter by capacity.disks.disk.name


|capacity.disks.usable_size
|integer
|query
|False
a|Filter by capacity.disks.usable_size


|capacity.disks.total_size
|integer
|query
|False
a|Filter by capacity.disks.total_size


|capacity.spare_allocation_units.syncmirror_pool
|string
|query
|False
a|Filter by capacity.spare_allocation_units.syncmirror_pool


|capacity.spare_allocation_units.size
|integer
|query
|False
a|Filter by capacity.spare_allocation_units.size


|capacity.spare_allocation_units.count
|integer
|query
|False
a|Filter by capacity.spare_allocation_units.count
```

```
|capacity.spare_allocation_units.node.name
|string
|query
|False
a|Filter by capacity.spare_allocation_units.node.name


|capacity.spare_allocation_units.node.uuid
|string
|query
|False
a|Filter by capacity.spare_allocation_units.node.uuid


|capacity.spare_allocation_units.available_size
|integer
|query
|False
a|Filter by capacity.spare_allocation_units.available_size


|capacity.disk_count
|integer
|query
|False
a|Filter by capacity.disk_count


|capacity.used_allocation_units.current_usage
|integer
|query
|False
a|Filter by capacity.used_allocation_units.current_usage


|capacity.used_allocation_units.node.name
|string
|query
|False
a|Filter by capacity.used_allocation_units.node.name


|capacity.used_allocation_units.node.uuid
|string
|query
|False
```

```
a|Filter by capacity.used_allocation_units.node.uuid


|capacity.used_allocation_units.count
|integer
|query
|False
a|Filter by capacity.used_allocation_units.count


|capacity.used_allocation_units.aggregate.name
|string
|query
|False
a|Filter by capacity.used_allocation_units.aggregate.name


|capacity.used_allocation_units.aggregate.uuid
|string
|query
|False
a|Filter by capacity.used_allocation_units.aggregate.uuid


|capacity.total
|integer
|query
|False
a|Filter by capacity.total


|capacity.remaining
|integer
|query
|False
a|Filter by capacity.remaining


|uuid
|string
|query
|False
a|Filter by uuid


|nodes.name
|string
```

```
|query
|False
a|Filter by nodes.name


|nodes.uuid
|string
|query
|False
a|Filter by nodes.uuid


|health.is_healthy
|boolean
|query
|False
a|Filter by health.is_healthy


|health.unhealthy_reason.arguments.message
|string
|query
|False
a|Filter by health.unhealthy_reason.arguments.message


|health.unhealthy_reason.arguments.code
|string
|query
|False
a|Filter by health.unhealthy_reason.arguments.code


|health.unhealthy_reason.code
|string
|query
|False
a|Filter by health.unhealthy_reason.code


|health.unhealthy_reason.message
|string
|query
|False
a|Filter by health.unhealthy_reason.message
```

```
|health.state
|string
|query
|False
a|Filter by health.state


|storage_type
|string
|query
|False
a|Filter by storage_type


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of shared storage pools in the cluster.


|records
|array[link:#storage_pool[storage_pool]]
a|

|spares
|array[link:#storage_pool_show_spares[storage_pool_show_spares]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "capacity": {
        "disks": [
          {
            "disk": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "1.0.1"
            },
            "total_size": 0,
            "usable_size": 0
```

```
          }
        ],
        "remaining": 0,
        "spare_allocation_units": [
          {
            "available_size": 0,
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "size": 0,
            "syncmirror_pool": "string"
          }
        ],
        "total": 0,
        "used_allocation_units": [
          {
            "aggregate": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "aggr1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "current_usage": 0,
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      },
      "health": {
        "state": "string",
```

```
        "unhealthy_reason": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        }
      },
      "name": "string",
      "nodes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "storage_type": "string",
      "uuid": "string"
    }
  ],
  "spares": [
    {
      "checksum_style": "string",
      "disk_class": "solid_state",
      "disk_type": "string",
      "nodes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "size": 10156769280,
      "syncmirror_pool": "string",
      "usable": 9
```

```
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 10944513
| Unable to retrieve shared disk capability information.

| 10944514
| Unable to enable shared disk capability.

| 10944527
| Storage pools are not supported in MetroCluster configurations.

| 10944528
| Unable to retrieve MetroCluster configuration information.

| 11206662
| There is no storage pool matching the specified UUID or name.

| 11206667
| Storage pool feature is not enabled.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#disk]
[.api-collapsible-fifth-title]
```

```
disk

Reference to the constituent disk object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===


[#storage_pool_disk]
[.api-collapsible-fifth-title]
storage_pool_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Reference to the constituent disk object.


|total_size
|integer
a|Raw capacity of the disk, in bytes.


|usable_size
|integer
a|Usable capacity of this disk, in bytes.
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

Specifies what node can use this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_spare_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_spare_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available_size
|integer
a|The usable capacity of this set of allocation units.


|count
```

```
|integer
a|The number of spare allocation units on this node.


|node
|link:#node[node]
a|Specifies what node can use this set of allocation units.


|size
|integer
a|Size of each allocation unit.


|syncmirror_pool
|string
a|The RAID SyncMirror Pool to which this allocation unit is assigned.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

The aggregate that is using this cache capacity.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

The node hosting the aggregate using this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_used_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_used_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|The aggregate that is using this cache capacity.


|count
```

```
|integer
a|The number of allocation units used by this aggregate.


|current_usage
|integer
a|The amount of cache space used by this aggregate.


|node
|link:#node[node]
a|The node hosting the aggregate using this set of allocation units.


|===


[#capacity]
[.api-collapsible-fifth-title]
capacity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|The number of disks in the storage pool.


|disks
|array[link:#storage_pool_disk[storage_pool_disk]]
a|Properties of each disk used in the shared storage pool.


|remaining
|integer
a|Remaining usable capacity in the flash pool, in bytes.


|spare_allocation_units
|array[link:#storage_pool_spare_allocation_unit[storage_pool_spare_allocat
ion_unit]]
a|Properties of spare allocation units.
```

```
|total
|integer
a|Total size of the flash pool, in bytes.


|used_allocation_units
|array[link:#storage_pool_used_allocation_unit[storage_pool_used_allocatio
n_unit]]
a|Information about the storage pool allocation units participating in the
cache tier of an aggregate.


|===


[#error]
[.api-collapsible-fifth-title]
error

Indicates why the storage pool is unhealthy. This property is not returned
for healthy storage pools.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===
```

```
[#health]
[.api-collapsible-fifth-title]
health

Properties that outline shared storage pool health.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|is_healthy
|boolean
a|Indicates whether the storage pool is able to participate in
provisioning operations.


|state
|string
a|The state of the shared storage pool.


|unhealthy_reason
|link:#error[error]
a|Indicates why the storage pool is unhealthy. This property is not
returned for healthy storage pools.


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|capacity
|link:#capacity[capacity]
a|

|health
|link:#health[health]
a|Properties that outline shared storage pool health.


|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
|string
```

```
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.


|===


[#storage_pool_show_spares]
[.api-collapsible-fifth-title]
storage_pool_show_spares

Available spares for stoarge pool.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|checksum_style
|string
a|The checksum type that has been assigned to the spares.


|disk_class
|string
a|Disk class of spares.


|disk_type
|string
a|Type of disk.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use the usable spares for storage pool.


|size
|integer
a|Usable size of each spare, in bytes.
```

```
|syncmirror_pool
|string
a|SyncMirror spare pool.


|usable
|integer
a|Total number of usable spares in the bucket. The usable count for each
class of spares does not include reserved spare capacity recommended by
ONTAP best practices.

* example: 9
* readOnly: 1
* Introduced in: 9.12
* x-nullable: true


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDe124110d8e2daf0837c2157dfa90fbe6]]
= Create a new storage pool

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/pools`#

*Introduced In:* 9.11

Creates a new storage pool using available solid state capacity attached
to the nodes specified.

== Required properties

The following properties are required in the POST body:

* `name` - Name of the new storage pool.
* `nodes[].name` or `nodes[].uuid` - Nodes that can use cache capacity
from the new storage pool. Only nodes in the same HA pair can be specified
for a given storage pool. Spare cache capacity will be distributed evenly
among the specified nodes.
* `capacity.disk_count` - Number of SSDs to be used to create the storage
pool.

== Related ONTAP commands

* `storage pool create`

== Example:

----
POST /api/storage/pools {"nodes": [{"name": "node1"}, {"name": "node2"}],
"name": "storage_pool_1", "capacity": {"disk_count": "4"}}
----
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|disk_size
|integer
|query
|False
a|If set, POST only selects SSDs within five percent of the specified
size.


|simulate
|boolean
|query
|False
a|When set to "true", the end point returns a simulated layout of the
proposed new storage pool, without changing system state.

* Introduced in: 9.12


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0
```

```
|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity
|link:#capacity[capacity]
a|

|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
|string
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.


|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "capacity": {
    "disks": [
      {
        "disk": {
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      }
    ],
    "remaining": 0,
    "spare_allocation_units": [
      {
        "available_size": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "string"
      }
    ],
    "total": 0,
    "used_allocation_units": [
      {
        "aggregate": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "current_usage": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "name": "string",
  "nodes": [
    {
      "name": "node1",
```

```
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "storage_type": "string",
  "uuid": "string"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11206666
| Storage pool is unhealthy.

| 11208658
| A storage pool already uses the specified name.

| 11208660
| Disk does not exist.

| 11208661
| Disk is not a spare disk.

| 11208662
| Disk is not an SSD.

| 11208663
| Disk is reserved for core dump.

| 11208664
| Could not determine checksum type of disk.

| 11208666
| Could not determine usable size of disk.

| 11208668
| Could not determine connectivity between controller and disk.

| 11208670
| Could not determine original owner of disk.

| 11208671
| Could not determine SyncMirror pool of disk.

| 11208673
| Could not determine HA mode of node.
```

```
| 11208674
| Could not determine HA partner of node.

| 11208675
| Disks specified in the disk list are not visible to node.

| 11208678
| The disk list contains disks from nodes which are not in HA pair.

| 11208679
| Storage pools are not supported on nodes.

| 11208680
| Internal error. Cannot determine configuration for node.

| 11208681
| Node is not online.

| 11208682
| Internal error. Sharing configuration mismatch.

| 11208684
| Unable to share disk.

| 11208686
| Disk cannot be shared.

| 11208687
| Unable to retrieve expected sharing configuration.

| 11208688
| Storage pool create job failed.

| 11208690
| Not all nodes sharing the storage pool view disk as a shared disk.

| 11208691
| Not enough matching spares available.

| 11208692
| A disk list or count is a required parameter for storage pool creation.

| 11208693
| Invalid number of disks specified.

| 11208698
| Internal error. Missing node name.
```

```
| 11208699
| Internal error. Missing partner name for node configured for HA.

| 11208701
| Node is a standalone node. Do not specify other nodes with a standalone
node.

| 11208703
| Incorrect number of nodes specified. Specify one node or both nodes in
an HA pair.

| 11208704
| Specified nodes are not part of HA relationship. Specify one node or
both nodes in an HA pair.

| 11208705
| Disk is a data center SSD, which cannot be used in storage pools.

| 11208706
| Disk is a SSD-ZNS, which cannot be used in storage pools.

| 11215756
| Missing a required field for POST request.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#disk]
[.api-collapsible-fifth-title]
disk

Reference to the constituent disk object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#storage_pool_disk]
[.api-collapsible-fifth-title]
storage_pool_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Reference to the constituent disk object.


|total_size
```

```
|integer
a|Raw capacity of the disk, in bytes.


|usable_size
|integer
a|Usable capacity of this disk, in bytes.


|===


[#node]
[.api-collapsible-fifth-title]
node

Specifies what node can use this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_spare_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_spare_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available_size
```

```
|integer
a|The usable capacity of this set of allocation units.


|count
|integer
a|The number of spare allocation units on this node.


|node
|link:#node[node]
a|Specifies what node can use this set of allocation units.


|size
|integer
a|Size of each allocation unit.


|syncmirror_pool
|string
a|The RAID SyncMirror Pool to which this allocation unit is assigned.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

The aggregate that is using this cache capacity.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

The node hosting the aggregate using this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_used_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_used_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|The aggregate that is using this cache capacity.


|count
|integer
a|The number of allocation units used by this aggregate.
```

```
|current_usage
|integer
a|The amount of cache space used by this aggregate.


|node
|link:#node[node]
a|The node hosting the aggregate using this set of allocation units.


|===


[#capacity]
[.api-collapsible-fifth-title]
capacity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|The number of disks in the storage pool.


|disks
|array[link:#storage_pool_disk[storage_pool_disk]]
a|Properties of each disk used in the shared storage pool.


|remaining
|integer
a|Remaining usable capacity in the flash pool, in bytes.


|spare_allocation_units
|array[link:#storage_pool_spare_allocation_unit[storage_pool_spare_allocat
ion_unit]]
a|Properties of spare allocation units.


|total
|integer
```

```
a|Total size of the flash pool, in bytes.


|used_allocation_units
|array[link:#storage_pool_used_allocation_unit[storage_pool_used_allocatio
n_unit]]
a|Information about the storage pool allocation units participating in the
cache tier of an aggregate.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Indicates why the storage pool is unhealthy. This property is not returned
for healthy storage pools.


[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#health]
[.api-collapsible-fifth-title]
health

Properties that outline shared storage pool health.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|is_healthy
|boolean
a|Indicates whether the storage pool is able to participate in
provisioning operations.


|state
|string
a|The state of the shared storage pool.


|===
```

```
[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity
|link:#capacity[capacity]
a|

|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
```

```
|string
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage specific storage pools

:leveloffset: +1


[[IDe90c151284d699d70b912db617ac28ae]]
= Manage specific storage pools



== Updating storage pools

The PATCH operation is used to modify properties of the storage pool.
There are several properties that can be modified on a storage pool.
PATCH operations on a storage pool are restricted when another PATCH
operation is in progress.
The following is a list of properties that can be modified using the PATCH
operation including a brief description for each:

* name - Can be updated to rename the storage pool.
* capacity.disk_count - Can be updated to increase the number of disks in
a storage pool.
```

```
* capacity.spare_allocation_units[].count - Modifying this value requires
that the user specify capacity.spare_allocation_units[].node as well.
Modifying this value redistributes spare cache capacity among the nodes
specified in the operation
When expanding a storage pool, the cache tiers of all aggregates using the
storage pool's allocation units are expanded automatically.

=== Simulated storage pool expansion

The PATCH operation also supports simulated expansion of a storage pool.
Running PATCH with the query "simulate" set to "true", and
"capacity.disk_count" set to the final disk count will return a response
containing the projected new capacity and the new constituent disk list
for the storage pool.

== Deleting storage pools

If cache capacity from a storage pool is being used in an aggregate, it
cannot be deleted.
See the /storage/aggregates API for details on deleting aggregates.

'''

== Examples

=== Retrieving a specific pool from the cluster

The following example shows the response of the requested storage pool. If
there is no storage pool with the requested UUID, an error is returned.

----

# The API:
/api/storage/pools/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/870dd9f2-bdfa-4167-b692-
57d1cec874d4" -H "accept: application/json"

# The response:
{
"uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
"name": "new_sp",
"storage_type": "ssd",
"nodes": [
  {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
```

```
    "name": "node-1",
  },
  {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  }
],
"capacity": {
    "remaining": 1846542336,
    "total": 7386169344,
    "spare_allocation_units": [
    {
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "count": 1,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 1846542336
    },
    {
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "count": 0,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 0
    }
    ],
    "used_allocation_units": [
    {
      "aggregate": {
        "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
        "name": "test_a"
      },
      "allocated_unit_count": 2,
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "capacity": 2769813504
    },
    {
```

```
      "aggregate": {
        "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
        "name": "test_b"
      },
      "allocated_unit_count": 1,
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "capacity": 1384906752
    }
    ],
    "disk_count": 4,
    "disks": [
      {
        "disk": {
          "name": "VMw-1.11"
        },
        "usable_size": 1902379008,
        "total_size": 1908871168,
      },
      {
        "disk": {
          "name": "VMw-1.12"
        },
        "usable_size": 1902379008,
        "total_size": 1908871168,
      },
      {
        "disk": {
          "name": "VMw-1.23"
        },
        "usable_size": 1902379008,
        "total_size": 1908871168,
      },
      {
        "disk": {
          "name": "VMw-1.24"
        },
        "usable_size": 1902379008,
        "total_size": 1908871168,
      }
    ]
  },
  "health": {
    "state": "normal",
```

```
    "is_healthy": true
  },
}
----

=== Simulating storage pool expansion

The following example shows the response for a simulated storage pool
expansion based on the values of the 'capacity.disk_count' attribute
passed in.
The query does not modify the existing storage pool, but rather returns
how it will look after the expansion.
This will be reflected in the following attributes:

* capacity.total- Total space, in bytes.
* capacity.remaining - New remaining capacity, in bytes.
* capacity.disks.disk - New list of constituent disks.
* capacity.disk_count - New number of disks in the pool.

----

# The API:
/api/storage/pools/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/cae60cfe-deae-42bd-
babb-ef437d118314?simulate=true" -H "accept: application/json" -d
"{\"capacity\": {\"disk_count\": 6}}"

# The response:
{
"records": [
{
  "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
  "name": "new_sp",
  "capacity": {
      "remaining": 1846542336,
      "total": 7386169344,
      "used_allocation_units": [
      {
        "aggregate": {
          "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
          "name": "test_a"
        },
        "current_usage": 2769813504
      },
```

```
      {
        "aggregate": {
          "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
          "name": "test_b"
        },
        "current_usage": 1384906752
      }
      ],
      "disk_count": 6,
      "disks": [
        {
          "disk": {
            "name": "VMw-1.11"
          },
        },
        {
          "disk": {
            "name": "VMw-1.12"
          },
        },
        {
          "disk": {
            "name": "VMw-1.23"
          },
        },
        {
          "disk": {
            "name": "VMw-1.24"
          },
        },
        {
          "disk": {
            "name": "VMw-1.13"
          },
        },
        {
          "disk": {
            "name": "VMw-1.14"
          },
        }
      ]
    }
}
]
}
----
```

```
=== Adding capacity to a storage pool

The following example shows the workflow of adding disks to the storage
pool.

Step 1: Check the current disk count on the storage pool.

----

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-
712f626c983c?fields=capacity.disk_count" -H "accept: application/json"

# The response:{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "sp1",
"capacity": {
  "disk_count": 4
}
}
----

Step 2: Update the pool with the new disk count in 'capacity.disk_count'.
The response to PATCH is a job unless the request is invalid.

----

# The API:
/api/storage/pools

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-
8f01-712f626c983c" -H "accept: application/hal+json" -d "{\"capacity\":
{\"disk_count\": 6}}"

# The response:
{
"job": {
  "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
    }
```

```
  }
}
}
----

Step 3: Wait for the job to finish, then call GET to see the reflected
change.

----

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-
712f626c983c?fields=capacity.disk_count" -H "accept: application/json"

# The response:
{
"uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
"name": "sp1",
"capacity": {
  "disk_count": 6
}
}
----

The following example shows the workflow to redistribute spare capacity
among nodes sharing the storage pool
Step 1: Check the current spare capacity distribution of the pool.

----

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-
5a402bffbe4b?fields=capacity.spare_allocation_units" -H "accept:
application/json"

# The response:
{
"uuid": "f3aafdc6-be35-4d93-9590-5a402bffbe4b",
"name": "sp1",
"capacity": {
    "spare_allocation_units": [
```

```
    {
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "count": 1,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 1846542336
    },
    {
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "count": 0,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 0
    }
    ],
}
}
----

Step 2: Update the pool so that the spare allocation unit count is
symmetrically modified for each node. The response to PATCH is a job
unless the request is invalid.

----

# The API:
/api/storage/pools

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-
9590-5a402bffbe4b" -H "accept: application/hal+json" -d '{ "capacity" : {
"spare_allocation_units": [{"node": {"name": "node-1"}, "count":0},
{"node": {"name": "node-2"}, "count": 1}]}}'

# The response:
{
"job": {
  "uuid": "6b7ab28e-168d-11ea-8a50-0050568eca76",
  "_links": {
    "self": {
```

```
      "href": "/api/cluster/jobs/6b7ab28e-168d-11ea-8a50-0050568eca76"
    }
  }
}
}
----

Step 3: Wait for the job to finish, then call GET to see the reflected
change.

----

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-
5a402bffbe4b?fields=capacity.spare_allocation_units" -H "accept:
application/json"

# The response:
{
"uuid": "f3aafdc6-be35-4d93-9590-5a402bffbe4b",
"name": "sp1",
"capacity": {
    "spare_allocation_units": [
    {
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "count": 0,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 0
    },
    {
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "count": 1,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 1846542336
    }
```

```
    ],
}
}
----




[[ID85f0eda4972dab921b5ed90b381c5e53]]
= Delete a storage pool specified by the UUID

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/pools/{uuid}`#

*Introduced In:* 9.11

Deletes the storage pool specified by the UUID. This request starts a job
and returns a link to that job.

== Related ONTAP commands

* `storage pool delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Storage Pool UUID.


|return_timeout
|integer
|query
|False
```

```
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11209668
| Capacity from pool is being used by one or more aggregates.
```

```
| 11209670
| Unable to stop sharing a constituent disk of the storage pool.

| 11209671
| Cannot find a node sharing this storage pool.

| 11209672
| Disk is not shared.

| 11209673
| Unable to assign shared capacity from a constituent disk.

| 11209675
| A constituent disk of the storage pool has non-spare capacity.

| 11209676
| A constituent disk of the storage pool is reserved for coredump.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
```

```
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
```

```
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID06f1da7e6717c7a64d14ecdeed703256]]
= Retrieve a storage pool specified by the UUID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/storage/pools/{uuid}`#

*Introduced In:* 9.11

Retrieves the storage pool specified by the UUID.

== Related ONTAP commands

* `storage pool show -storage-pool-uuid`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Storage pool UUID.


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|capacity
|link:#capacity[capacity]
a|

|health
|link:#health[health]
a|Properties that outline shared storage pool health.


|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
|string
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```
  "capacity": {
    "disks": [
      {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      }
    ],
    "remaining": 0,
    "spare_allocation_units": [
      {
        "available_size": 0,
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "string"
      }
    ],
    "total": 0,
    "used_allocation_units": [
      {
        "aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "current_usage": 0,
        "node": {
```

```
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "health": {
    "state": "string",
    "unhealthy_reason": {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist"
    }
  },
  "name": "string",
  "nodes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "storage_type": "string",
  "uuid": "string"
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11206662
| There is no storage pool matching the specified UUID or name.

| 11215856
| The specified storage pool is not healthy.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
```

```
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#disk]
[.api-collapsible-fifth-title]
disk
```

```
Reference to the constituent disk object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===


[#storage_pool_disk]
[.api-collapsible-fifth-title]
storage_pool_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Reference to the constituent disk object.


|total_size
|integer
a|Raw capacity of the disk, in bytes.


|usable_size
|integer
a|Usable capacity of this disk, in bytes.


|===
```

```
[#node]
[.api-collapsible-fifth-title]
node

Specifies what node can use this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_spare_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_spare_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available_size
|integer
a|The usable capacity of this set of allocation units.


|count
|integer
```

```
a|The number of spare allocation units on this node.


|node
|link:#node[node]
a|Specifies what node can use this set of allocation units.


|size
|integer
a|Size of each allocation unit.


|syncmirror_pool
|string
a|The RAID SyncMirror Pool to which this allocation unit is assigned.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

The aggregate that is using this cache capacity.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===
```

```
[#node]
[.api-collapsible-fifth-title]
node

The node hosting the aggregate using this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_used_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_used_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|The aggregate that is using this cache capacity.


|count
|integer
```

```
a|The number of allocation units used by this aggregate.


|current_usage
|integer
a|The amount of cache space used by this aggregate.


|node
|link:#node[node]
a|The node hosting the aggregate using this set of allocation units.


|===


[#capacity]
[.api-collapsible-fifth-title]
capacity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|The number of disks in the storage pool.


|disks
|array[link:#storage_pool_disk[storage_pool_disk]]
a|Properties of each disk used in the shared storage pool.


|remaining
|integer
a|Remaining usable capacity in the flash pool, in bytes.


|spare_allocation_units
|array[link:#storage_pool_spare_allocation_unit[storage_pool_spare_allocat
ion_unit]]
a|Properties of spare allocation units.
```

```
|total
|integer
a|Total size of the flash pool, in bytes.


|used_allocation_units
|array[link:#storage_pool_used_allocation_unit[storage_pool_used_allocatio
n_unit]]
a|Information about the storage pool allocation units participating in the
cache tier of an aggregate.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Indicates why the storage pool is unhealthy. This property is not returned
for healthy storage pools.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#health]
[.api-collapsible-fifth-title]
health

Properties that outline shared storage pool health.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|is_healthy
|boolean
a|Indicates whether the storage pool is able to participate in
provisioning operations.


|state
|string
a|The state of the shared storage pool.
```

```
|unhealthy_reason
|link:#error[error]
a|Indicates why the storage pool is unhealthy. This property is not
returned for healthy storage pools.


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID655383cdb9ee0269612088ec8c3bb9c9]]
= Update a storage pool specified by the UUID

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/pools/{uuid}`#

*Introduced In:* 9.11

Updates the storage pool specified by the UUID with the properties in the
body. This request starts a job and returns a link to that job.

== Related ONTAP commands

* `storage pool rename`
* `storage pool reassign`
* `storage pool add`


== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Storage pool UUID.


|simulate
|boolean
|query
|False
a|When set to "true" in conjunction with a PATCH on "capacity.disk_count",
the end point returns a simulated layout of the storage pool with the
additional disks, without changing system state.

* Introduced in: 9.12


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity
|link:#capacity[capacity]
a|

|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
|string
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "capacity": {
    "disks": [
      {
        "disk": {
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      }
```

```
    ],
    "remaining": 0,
    "spare_allocation_units": [
      {
        "available_size": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "string"
      }
    ],
    "total": 0,
    "used_allocation_units": [
      {
        "aggregate": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "current_usage": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "name": "string",
  "nodes": [
    {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "storage_type": "string",
  "uuid": "string"
}
====

== Response
```


Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|num_records
|integer
a|Number of shared storage pools in the cluster.


|records
|array[link:#storage_pool[storage_pool]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  },
  "num_records": 1,
  "records": [
    {
      "capacity": {
        "disks": [
          {
            "disk": {
              "name": "1.0.1"
            },
            "total_size": 0,
            "usable_size": 0
          }
        ],
        "remaining": 0,
        "spare_allocation_units": [
          {
            "available_size": 0,
```

```
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "size": 0,
            "syncmirror_pool": "string"
          }
        ],
        "total": 0,
        "used_allocation_units": [
          {
            "aggregate": {
              "name": "aggr1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "current_usage": 0,
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      },
      "name": "string",
      "nodes": [
        {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "storage_type": "string",
      "uuid": "string"
    }
  ]
}
====
```


```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11211658
| Node does not have enough spare capacity.

| 11211659
| Valid allocation unit input is required.

| 11211662
| Specified node is not part of the storage pool.

| 11211663
| Failed to reassign available capacity in the storage pool.

| 11211664
| Could not fix the broken allocation unit for the storage pool.

| 11212673
| Could not grow one or more aggregates.

| 11212679
| Adding specified number of disks will expand storage pool beyond maximum
supported disk limit.

| 11212680
| Incorrect node specified.

| 11212681
| 0 is an invalid value for disk_count.

| 11212682
| Adding the specified number of disks will result in the storage pool
reaching the maximum disk limit reserved for RAID-TEC use only. At this
limit, the storage pool can only allocate capacity to aggregates
containing RAID-TEC RAID groups. Existing aggregates containing RAID
groups other than RAID-TEC will not automatically grow to the new
capacity.

| 11212683
| Renaming storage pool to new name failed.

| 11212763
| Storage pool add job failed.
```

```
| 11215657
| Storage pool PATCH request have missing parameters.

| 11215658
| Storage pool PATCH request for reassign is invalid.

| 11215659
| Storage pool PATCH request for reassign have invalid allocation unit
count.

| 11215660
| Storage pool PATCH request for reassign have invalid node name.

| 11215662
| Storage pool PATCH request have invalid disk count.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
```

```
[.api-collapsible-fifth-title]
_links
[#disk]
[.api-collapsible-fifth-title]
disk

Reference to the constituent disk object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#storage_pool_disk]
[.api-collapsible-fifth-title]
storage_pool_disk

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk
|link:#disk[disk]
a|Reference to the constituent disk object.


|total_size
|integer
a|Raw capacity of the disk, in bytes.


|usable_size
|integer
a|Usable capacity of this disk, in bytes.
```

```
|===


[#node]
[.api-collapsible-fifth-title]
node

Specifies what node can use this set of allocation units.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_spare_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_spare_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available_size
|integer
a|The usable capacity of this set of allocation units.


|count
|integer
a|The number of spare allocation units on this node.
```

```
|node
|link:#node[node]
a|Specifies what node can use this set of allocation units.


|size
|integer
a|Size of each allocation unit.


|syncmirror_pool
|string
a|The RAID SyncMirror Pool to which this allocation unit is assigned.


|===


[#aggregate]
[.api-collapsible-fifth-title]
aggregate

The aggregate that is using this cache capacity.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

The node hosting the aggregate using this set of allocation units.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool_used_allocation_unit]
[.api-collapsible-fifth-title]
storage_pool_used_allocation_unit

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|The aggregate that is using this cache capacity.


|count
|integer
a|The number of allocation units used by this aggregate.


|current_usage
|integer
a|The amount of cache space used by this aggregate.


|node
|link:#node[node]
a|The node hosting the aggregate using this set of allocation units.
```

```
|===


[#capacity]
[.api-collapsible-fifth-title]
capacity

[cols=3*,options=header]
|===
|Name
|Type
|Description

|disk_count
|integer
a|The number of disks in the storage pool.


|disks
|array[link:#storage_pool_disk[storage_pool_disk]]
a|Properties of each disk used in the shared storage pool.


|remaining
|integer
a|Remaining usable capacity in the flash pool, in bytes.


|spare_allocation_units
|array[link:#storage_pool_spare_allocation_unit[storage_pool_spare_allocat
ion_unit]]
a|Properties of spare allocation units.


|total
|integer
a|Total size of the flash pool, in bytes.


|used_allocation_units
|array[link:#storage_pool_used_allocation_unit[storage_pool_used_allocatio
n_unit]]
a|Information about the storage pool allocation units participating in the
cache tier of an aggregate.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

Indicates why the storage pool is unhealthy. This property is not returned
for healthy storage pools.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
```

```
|string
a|Error code


|message
|string
a|Error message


|===


[#health]
[.api-collapsible-fifth-title]
health

Properties that outline shared storage pool health.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|is_healthy
|boolean
a|Indicates whether the storage pool is able to participate in
provisioning operations.


|state
|string
a|The state of the shared storage pool.


|===


[#node_reference]
[.api-collapsible-fifth-title]
node_reference

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_pool]
[.api-collapsible-fifth-title]
storage_pool

[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity
|link:#capacity[capacity]
a|

|name
|string
a|Storage pool name.


|nodes
|array[link:#node_reference[node_reference]]
a|Nodes that can use this storage pool for their aggregates.


|storage_type
|string
a|Storage type for the disks used to create the storage pool.


|uuid
|string
a|Storage pool UUID.
```

```
|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage ports

:leveloffset: +1


[[ID8eaf11d3de58654a4c8c81db806f3242]]
= Manage storage ports



== Retrieving storage port information

The storage port GET API retrieves all of the storage ports in the
cluster.

'''

== Examples

=== 1) Retrieve a list of storage ports from the cluster

==== The following example shows the response with a list of storage ports
in the cluster:

'''

----

# The API:
/api/storage/ports
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/storage/ports" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0a",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0a"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0b",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0b"
      }
    }
  },
  {
```

```
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0c",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0c"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0d",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0d"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
```

```
      }
    },
    "name": "0e",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0e"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0f",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0f"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0g",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0g"
      }
```

```
    }
  },
],
"num_records": 7,
"_links": {
  "self": {
    "href": "/api/storage/ports"
  }
}
}
----

'''

=== 2) Retrieve a specific storage port from the cluster

==== The following example shows the response of the requested storage
port. If there is no storage port with the requested node uuid and name,
an error is returned.

'''

----

# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0a" -H "accept: application/hal+json"

# The response:
{
"node": {
  "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
  "name": "node-1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
    }
  }
},
"name": "0a",
"description": "SAS Host Adapter 0a (PMC-Sierra PM8001 rev. C)",
"wwn": "500a098003633df0",
"speed": 6,
```

```
"cable": {
  "part_number": "112-00429+A0",
  "serial_number": "629230774",
  "identifier": "500a0980066e2c01-500a098003633df0",
  "length": "0.5m",
  "vendor": "Molex Inc.",
  "transceiver": "mini_sas_hd"
},
"state": "online",
"enabled": true,
"firmware_version": "01.12.09.00",
"type": "sas",
"redundant": true,
"in_use": true,
"_links": {
  "self": {
    "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a"
  }
}
}
----

'''

== Updating a storage port

The storage port PATCH API modifies the port mode for storage/network use
and allows the port to be enabled/disabled.

'''

== Examples

=== 1) Using an Ethernet port for storage

The following example sets an Ethernet port mode for storage use:

'''

----

# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-
907f-00a0985a72ee/e3a" -H "accept: application/hal+json" -H "Content-Type:
```

```
application/hal_json" -d '{"mode": "storage"}'

# The response:
{
}
----

'''

=== 2) Disabling a storage port

The following example disables an unused storage port:

'''

----

# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-
907f-00a0985a72ee/e3a" -H "accept: application/hal+json" -H "Content-Type:
application/hal_json" -d '{"enabled": "false"}'

# The response:
{
}
----

'''




[[ID30f7f2bce51263c681d91674822816a4]]
= Retrieve storage ports

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/ports`#

*Introduced In:* 9.6

Retrieves a collection of storage ports.

== Related ONTAP commands
```

```
* `storage port show`

== Learn more

* xref:{relative_path}storage_ports_endpoint_overview.html[DOC
/storage/ports]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|mode
|string
|query
|False
a|Filter by mode

* Introduced in: 9.8


|in_use
|boolean
|query
|False
a|Filter by in_use

* Introduced in: 9.9


|redundant
|boolean
|query
|False
a|Filter by redundant

* Introduced in: 9.9
```

```
|name
|string
|query
|False
a|Filter by name


|state
|string
|query
|False
a|Filter by state


|description
|string
|query
|False
a|Filter by description


|mac_address
|string
|query
|False
a|Filter by mac_address


|speed
|number
|query
|False
a|Filter by speed


|wwpn
|string
|query
|False
a|Filter by wwpn

* Introduced in: 9.9


|enabled
|boolean
|query
```

```
|False
a|Filter by enabled

* Introduced in: 9.9


|node.name
|string
|query
|False
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|type
|string
|query
|False
a|Filter by type

* Introduced in: 9.9


|cable.part_number
|string
|query
|False
a|Filter by cable.part_number


|cable.transceiver
|string
|query
|False
a|Filter by cable.transceiver

* Introduced in: 9.15


|cable.length
|string
```

```
|query
|False
a|Filter by cable.length


|cable.serial_number
|string
|query
|False
a|Filter by cable.serial_number


|cable.identifier
|string
|query
|False
a|Filter by cable.identifier


|cable.vendor
|string
|query
|False
a|Filter by cable.vendor

* Introduced in: 9.15


|part_number
|string
|query
|False
a|Filter by part_number


|board_name
|string
|query
|False
a|Filter by board_name


|serial_number
|string
|query
|False
a|Filter by serial_number
```

```
|error.message
|string
|query
|False
a|Filter by error.message


|error.corrective_action
|string
|query
|False
a|Filter by error.corrective_action


|firmware_version
|string
|query
|False
a|Filter by firmware_version

* Introduced in: 9.9


|wwn
|string
|query
|False
a|Filter by wwn


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
```

```
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#storage_port[storage_port]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "board_name": "string",
      "cable": {
        "identifier": "500a0980000b6c3f-50000d1703544b80",
        "length": "2m",
        "part_number": "112-00431+A0",
        "serial_number": 616930439,
        "transceiver": "mini_sas_hd",
        "vendor": "Molex Inc."
      },
      "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
      "error": {
        "corrective_action": "string",
        "message": "string"
      },
      "firmware_version": "03.08.09.00",
      "mac_address": "string",
      "mode": "storage",
      "name": "2a",
```

```
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "part_number": "111-03801",
      "serial_number": "7A2463CC45B",
      "speed": 6,
      "state": "online",
      "type": "sas",
      "wwn": "50000d1703544b80",
      "wwpn": "string"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
```

```
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
a|

|length
|string
a|

|part_number
|string
a|

|serial_number
|string
a|

|transceiver
|string
a|

|vendor
|string
a|

|===
```

```
[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|corrective_action
|string
a|Error corrective action


|message
|string
a|Error message


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#storage_port]
[.api-collapsible-fifth-title]
storage_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|board_name
|string
a|

|cable
|link:#cable[cable]
a|

|description
|string
a|

|enabled
|boolean
a|

|error
|link:#error[error]
```

```
a|

|firmware_version
|string
a|

|force
|boolean
a|

|in_use
|boolean
a|Specifies whether any devices are connected through this port


|mac_address
|string
a|

|mode
|string
a|Operational mode of a non-dedicated Ethernet port


|name
|string
a|

|node
|link:#node[node]
a|

|part_number
|string
a|

|redundant
|boolean
a|Specifies whether all devices connected through this port have a
redundant path from another port


|serial_number
|string
a|

|speed
```

```
|number
a|Operational port speed in Gbps


|state
|string
a|

|type
|string
a|

|wwn
|string
a|World Wide Name


|wwpn
|string
a|World Wide Port Name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd3bea53c1b05f5850b6eaaf5f7c9b7cc]]
= Retrieve a storage port

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/ports/{node.uuid}/{name}`#

*Introduced In:* 9.6
```

```
Retrieves a specific storage port.

== Related ONTAP commands

* `storage port show`

== Learn more

* xref:{relative_path}storage_ports_endpoint_overview.html[DOC
/storage/ports]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|name
|string
|path
|True
a|Port name


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|board_name
|string
a|

|cable
|link:#cable[cable]
a|

|description
|string
a|

|enabled
|boolean
a|

|error
|link:#error[error]
a|

|firmware_version
|string
a|

|force
|boolean
a|

|in_use
|boolean
a|Specifies whether any devices are connected through this port


|mac_address
|string
a|

|mode
```

```
|string
a|Operational mode of a non-dedicated Ethernet port


|name
|string
a|

|node
|link:#node[node]
a|

|part_number
|string
a|

|redundant
|boolean
a|Specifies whether all devices connected through this port have a
redundant path from another port


|serial_number
|string
a|

|speed
|number
a|Operational port speed in Gbps


|state
|string
a|

|type
|string
a|

|wwn
|string
a|World Wide Name


|wwpn
|string
a|World Wide Port Name
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "board_name": "string",
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439,
    "transceiver": "mini_sas_hd",
    "vendor": "Molex Inc."
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "error": {
    "corrective_action": "string",
    "message": "string"
  },
  "firmware_version": "03.08.09.00",
  "mac_address": "string",
  "mode": "storage",
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": 6,
  "state": "online",
  "type": "sas",
  "wwn": "50000d1703544b80",
  "wwpn": "string"
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#cable]
[.api-collapsible-fifth-title]
```

```
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
a|

|length
|string
a|

|part_number
|string
a|

|serial_number
|string
a|

|transceiver
|string
a|

|vendor
|string
a|

|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|corrective_action
|string
```

```
a|Error corrective action


|message
|string
a|Error message


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
```

```
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0e6d503b39449e591b27c2402bcda0da]]
= Update a storage port

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/ports/{node.uuid}/{name}`#

*Introduced In:* 9.11

Updates a storage port.
```

```
== Related ONTAP commands

* `storage port modify`
* `storage port enable`
* `storage port disable`

== Learn more

* xref:{relative_path}storage_ports_endpoint_overview.html[DOC
/storage/ports]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|name
|string
|path
|True
a|Port name

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|board_name
|string
a|

|cable
|link:#cable[cable]
a|

|description
|string
a|

|enabled
|boolean
a|

|error
|link:#error[error]
a|

|firmware_version
|string
a|

|force
|boolean
a|

|in_use
|boolean
a|Specifies whether any devices are connected through this port


|mac_address
|string
a|

|mode
|string
a|Operational mode of a non-dedicated Ethernet port


|name
|string
a|
```

```
|node
|link:#node[node]
a|

|part_number
|string
a|

|redundant
|boolean
a|Specifies whether all devices connected through this port have a
redundant path from another port


|serial_number
|string
a|

|speed
|number
a|Operational port speed in Gbps


|state
|string
a|

|type
|string
a|

|wwn
|string
a|World Wide Name


|wwpn
|string
a|World Wide Port Name


|===


.Example request
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "board_name": "string",
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439,
    "transceiver": "mini_sas_hd",
    "vendor": "Molex Inc."
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "error": {
    "corrective_action": "string",
    "message": "string"
  },
  "firmware_version": "03.08.09.00",
  "mac_address": "string",
  "mode": "storage",
  "name": "2a",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": 6,
  "state": "online",
  "type": "sas",
  "wwn": "50000d1703544b80",
  "wwpn": "string"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 17891328
| Port operation "+++<operation>+++\" failed on port \"+++<name>+++\".
This might indicate a hardware error, an illegal request, or an aborted
command.+++</name>++++++</operation>+++

| 17891329
| Port "+++<name>+++\" is invalid.+++</name>+++

| 17891330
| Port operation "+++<operation>+++\" failed on port \"+++<port>+++\"
because it is not supported on this port
type.+++</port>++++++</operation>+++

| 17891331
| Cannot complete operation on port "+++<port>+++\". The status of the
port is not available.+++</port>+++

| 17891332
| Port operation "+++<operation>+++\" failed on port \"+++<port>+++\". The
port is not offline.+++</port>++++++</operation>+++

| 17891333
| Port operation "+++<operation>+++\" failed on port \"+++<port>+++\". The
port is already offline.+++</port>++++++</operation>+++

| 17891334
| Port operation "+++<operation>+++\" failed on port \"+++<port>+++\". One
or more devices on the port is currently being
sanitized.+++</port>++++++</operation>+++

| 17891338
| Device operation "+++<operation>+++\" failed on port \"+++<name>+++\".
Some devices can only be accessed through this port. Disabling this port
might cause multiple device failures.+++</name>++++++</operation>+++

| 17891339
| Device operation "+++<operation>+++\" failed on port \"+++<name>+++\".
Some devices can only be accessed through this port. Disabling this port
might cause multiple device failures. Use the \"-force\" parameter to
ignore checks and +++<operation>+++the
port.+++</operation>++++++</name>++++++</operation>+++

| 17891344
| Device operation "+++<operation>+++\" failed on port \"+++<name>+++\"
because the port is in use.+++</name>++++++</operation>+++
```

```
| 17891345
| Device operation "+++<operation>+++\" failed on port \"+++<name>+++\"
because the port is in use. Use the \"-force\" parameter to ignore checks
and +++<operation>+++the
port.+++</operation>++++++</name>++++++</operation>+++

| 17891346
| Port "+++<name>+++\" is already enabled.+++</name>+++

| 17891347
| Port "+++<name>+++\" is already disabled.+++</name>+++

| 17891350
| Port "+++<name>+++\" is not a valid storage Ethernet port.+++</name>+++

| 17891350
| Port "+++<port>+++\" is not a valid storage Ethernet port.+++</port>+++

| 17891352
| Port operation "+++<operation>+++\" failed on port \"+++<name>+++\".
Reboot node \"+++<node>+++\" to
recover.+++</node>++++++</name>++++++</operation>+++

| 17891353
| Unable to enable port "+++<port>+++\" because it is in network
mode.+++</port>+++

| 17891354
| Unable to disable port "+++<port>+++\" because it is in network
mode.+++</port>+++

| 17891355
| Port operation "+++<operation>+++\" failed on port \"+++<name>+++\"
because it is not supported on dedicated
ports.+++</name>++++++</operation>+++

| 17891356
| Unable to +++<operation>+++port \"+++<port>+++\" when setting mode to
\"+++<mode>+++\".+++</mode>++++++</port>++++++</operation>+++
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
a|

|length
|string
a|

|part_number
|string
a|

|serial_number
|string
a|

|transceiver
|string
a|

|vendor
|string
a|

|===


[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|corrective_action
|string
a|Error corrective action


|message
|string
a|Error message


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#storage_port]
[.api-collapsible-fifth-title]
storage_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|board_name
|string
a|

|cable
|link:#cable[cable]
a|

|description
|string
a|

|enabled
|boolean
a|

|error
|link:#error[error]
a|

|firmware_version
|string
a|
```

```
|force
|boolean
a|

|in_use
|boolean
a|Specifies whether any devices are connected through this port


|mac_address
|string
a|

|mode
|string
a|Operational mode of a non-dedicated Ethernet port


|name
|string
a|

|node
|link:#node[node]
a|

|part_number
|string
a|

|redundant
|boolean
a|Specifies whether all devices connected through this port have a
redundant path from another port


|serial_number
|string
a|

|speed
|number
a|Operational port speed in Gbps


|state
```

```
|string
a|

|type
|string
a|

|wwn
|string
a|World Wide Name


|wwpn
|string
a|World Wide Port Name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage QoS policies

:leveloffset: +1


[[ID170afd0e2f9013effc37ef9e830dd976]]
= Manage QoS policies
```

```
== Quality of Service Configuration

A QoS policy defines measurable service level objectives (SLOs) that apply
to the storage objects with which the policy is associated. There are two
types of policies that can be configured: fixed, which defines a fixed
SLO, or adaptive which defines a variable SLO for a storage object.
Adaptive policies vary the SLO depending on the space usage of the storage
object. A policy can be either a fixed policy or an adaptive one, not
both.
 +
Service level objectives include minimum and maximum limits on throughput
in terms of IOPS. Only maximum limits can be set in terms of both IOPS
and/or throughput (MB/s). A QoS policy can be used to enforce SLOs for
multiple storage objects by specifying "capacity_shared" to true. For
example, if a QoS policy with "capacity_shared" is set to true and it has
maximum_throughput_iops set to 1000, and this policy is assigned to four
volumes, then the combined throughput of all four volumes is limited to
1000 IOPS. If "capacity_shared" is set to false then, each storage object
will have it's SLOs enforced individually. For example, in the previous
case if the same policy was applied to four volumes but with
"capacity_shared" set to false, then each of the volumes would be limited
to 1000 IOPS individually. Once "capacity_shared" is set, it cannot be
modified.
 +
Adaptive parameters can specify the variable SLOs in terms of IOPS/TB. The
actual IOPS enforced on the storage object can be calculated using the
allocated space on the storage object. The policies are enforced
individually amongst storage objects.

== Examples

=== 1) Create a fixed QoS policy

The following example shows how to create a fixed QoS policy to limit
throughput for a storage object between 5000 IOPS and 10000 IOPS which has
capacity_shared set to false. This QoS policy can be used as a template to
apply on multiple storage objects to provide individual SLOs to each
object.
 +

'''

----
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
```

```
\"fixed\": { \"capacity_shared\": false, \"max_throughput_iops\": 10000,
\"min_throughput_iops\": 5000 }, \"name\":
\"qos_policy_5000_to_10000_iops\", \"svm\": { \"name\": \"vs0\" }}"
----

'''

=== 2) Create an adaptive QoS policy

The following example shows how to create an adaptive QoS policy which
provides 5000 IOPS per GB of allocated space for a storage object with a
peak of 6000 IOPS. Minimum IOPS regardless of allocated space are 1000
IOPS.
 +

'''

----
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"adaptive\": { \"absolute_min_iops\": 1000, \"expected_iops\": 5000,
\"expected_iops_allocation\": \"used_space\", \"peak_iops\": 6000,
\"peak_iops_allocation\": \"allocated_space\" }, \"name\":
\"adaptive_pg_5k_to_6k\", \"svm\": { \"name\": \"vs0\" }}"
----

-

'''

=== 3) Update an existing QoS policy

The following example shows how to update SLOs of an existing QoS policy
and also rename it.
 +

'''

----
curl -X PATCH "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fixed\": {
\"max_throughput_iops\": 15000, \"min_throughput_iops\": 10000 },
\"name\": \"qos_policy_10k_to_15k_iops\"}"
----
```

```
'''

=== 4) Delete an existing QoS policy

When a QoS policy is deleted any associations of the policy with a storage
objects are also removed.
 +

'''

----
curl -X DELETE "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json"
----

'''





[[ID d87cb222ce5afcc58d5e0902afbab2ab]]
= Retrieve QoS policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qos/policies`#

*Introduced In:* 9.6

Retrieves a collection of QoS policies.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|object_count
|integer
|query
```

```
|False
a|Filter by object_count


|fixed.max_throughput_iops
|integer
|query
|False
a|Filter by fixed.max_throughput_iops

* Max value: 2147483647
* Min value: 0


|fixed.min_throughput_mbps
|integer
|query
|False
a|Filter by fixed.min_throughput_mbps

* Introduced in: 9.8
* Max value: 4194303
* Min value: 0


|fixed.capacity_shared
|boolean
|query
|False
a|Filter by fixed.capacity_shared


|fixed.max_throughput_mbps
|integer
|query
|False
a|Filter by fixed.max_throughput_mbps

* Max value: 4194303
* Min value: 0


|fixed.min_throughput_iops
|integer
|query
|False
a|Filter by fixed.min_throughput_iops
```

```
* Max value: 2147483647
* Min value: 0


|fixed.max_throughput
|string
|query
|False
a|Filter by fixed.max_throughput

* Introduced in: 9.17


|fixed.min_throughput
|string
|query
|False
a|Filter by fixed.min_throughput

* Introduced in: 9.17


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|name
|string
|query
|False
a|Filter by name


|scope
|string
|query
```

```
|False
a|Filter by scope

* Introduced in: 9.11


|pgid
|integer
|query
|False
a|Filter by pgid

* Introduced in: 9.10


|uuid
|string
|query
|False
a|Filter by uuid


|policy_class
|string
|query
|False
a|Filter by policy_class

* Introduced in: 9.10


|adaptive.absolute_min_iops
|integer
|query
|False
a|Filter by adaptive.absolute_min_iops


|adaptive.expected_iops_allocation
|string
|query
|False
a|Filter by adaptive.expected_iops_allocation

* Introduced in: 9.10
```

```
|adaptive.block_size
|string
|query
|False
a|Filter by adaptive.block_size

* Introduced in: 9.10


|adaptive.expected_iops
|integer
|query
|False
a|Filter by adaptive.expected_iops


|adaptive.peak_iops
|integer
|query
|False
a|Filter by adaptive.peak_iops


|adaptive.peak_iops_allocation
|string
|query
|False
a|Filter by adaptive.peak_iops_allocation

* Introduced in: 9.10


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
```

```
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#qos_policy[qos_policy]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
      },
      "adaptive": {
        "block_size": "string",
        "expected_iops_allocation": "string",
        "peak_iops_allocation": "string"
      },
      "fixed": {
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ]
      },
      "name": "extreme",
      "object_count": 0,
      "pgid": 0,
      "policy_class": "string",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
```

```
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|self
|link:#href[href]
a|

|===


[#adaptive]
[.api-collapsible-fifth-title]
adaptive

Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|absolute_min_iops
|integer
a|Specifies the absolute minimum IOPS that is used as an override when the
expected_iops is less than this value. These floors are not guaranteed on
non-AFF platforms or when FabricPool tiering policies are set.


|block_size
|string
a|Specifies the block size


|expected_iops
|integer
a|Expected IOPS. Specifies the minimum expected IOPS per TB allocated
based on the storage object allocated size. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|expected_iops_allocation
|string
a|Specifies the size to be used to calculate expected IOPS per TB. The
```

```
size options are either the storage object allocated space or the storage
object used space.


|peak_iops
|integer
a|Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on
the storage object allocated size or the storage object used size.


|peak_iops_allocation
|string
a|Specifies the size to be used to calculate peak IOPS per TB. The size
options are either the storage object allocated space or the storage
object used space.


|===


[#fixed]
[.api-collapsible-fifth-title]
fixed

QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity_shared
|boolean
a|Specifies whether the capacities are shared across all objects that use
this QoS policy-group. Default is false.


|max_throughput
|string
a|Maximum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no maximum
throughput is enforced.
```

```
|max_throughput_iops
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no maximum throughput is enforced.


|max_throughput_mbps
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no maximum throughput is enforced.


|min_throughput
|string
a|Minimum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no minimum
throughput is enforced.


|min_throughput_iops
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no minimum throughput is enforced. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|min_throughput_mbps
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no minimum throughput is enforced.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|adaptive
|link:#adaptive[adaptive]
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.
```

```
|name
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID94edbb81997e79142f96d0568f582e5b]]
= Create a QoS policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/qos/policies`#

*Introduced In:* 9.6

Creates a QoS policy.

== Required properties

* `svm.uuid` or `svm.name` - The existing SVM owning the QoS policy.
* `name` - The name of the QoS policy.
* `fixed.+*+` or `adaptive.+*+` - Either of the fixed or adaptive
parameters.

== Default property values
```

```
* If `fixed.+*+` parameters are specified, then capacity.shared is set to
false by default.

== Related ONTAP commands

* `qos policy-group create`
* `qos adaptive-policy-group create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|adaptive
|link:#adaptive[adaptive]
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


|name
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
```

```
"svm" for SVM owned objects.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "adaptive": {
    "block_size": "string",
    "expected_iops_allocation": "string",
    "peak_iops_allocation": "string"
  },
  "fixed": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ]
  },
  "name": "extreme",
```

```
  "object_count": 0,
  "pgid": 0,
  "policy_class": "string",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 8454147
| The maximum limit for QoS policies has been reached.

| 8454154
| The name specified for creating conflicts with an existing QoS policy
name.

| 8454194
| The minimum throughput value for the policy group must be less than or
equal to the maximum throughput value.

| 8454260
| Invalid value for maximum and minimum fields. Valid values for
max_throughput_iops and max_throughput_mbps combination is for the ratio
of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.

| 8454273
| Invalid value for an adaptive field. Value should be non-zero.

| 8454274
| Invalid value for an adaptive field. Value for expected_iops must be
less than or equal to the value for peak_iops.

| 8454277
| The name specified for creating an adaptive QoS policy conflicts with an
existing fixed QoS policy name.

| 8454278
| The name specified for creating a fixed QoS policy conflicts with an
existing adaptive QoS policy name.

| 8454379
| The name specified for creating a fixed QoS policy already exists.

| 8454380
| The name specified for creating an adaptive QoS policy already exists.
```

```
| 8454389
| One specified field cannot be used in combination with at least one
other specified field.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#adaptive]
[.api-collapsible-fifth-title]
adaptive

Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|absolute_min_iops
|integer
a|Specifies the absolute minimum IOPS that is used as an override when the
expected_iops is less than this value. These floors are not guaranteed on
non-AFF platforms or when FabricPool tiering policies are set.


|block_size
|string
a|Specifies the block size


|expected_iops
|integer
a|Expected IOPS. Specifies the minimum expected IOPS per TB allocated
based on the storage object allocated size. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|expected_iops_allocation
|string
a|Specifies the size to be used to calculate expected IOPS per TB. The
size options are either the storage object allocated space or the storage
object used space.


|peak_iops
|integer
a|Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on
the storage object allocated size or the storage object used size.


|peak_iops_allocation
|string
a|Specifies the size to be used to calculate peak IOPS per TB. The size
options are either the storage object allocated space or the storage
object used space.


|===
```

```
[#fixed]
[.api-collapsible-fifth-title]
fixed

QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity_shared
|boolean
a|Specifies whether the capacities are shared across all objects that use
this QoS policy-group. Default is false.


|max_throughput
|string
a|Maximum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no maximum
throughput is enforced.


|max_throughput_iops
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no maximum throughput is enforced.


|max_throughput_mbps
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no maximum throughput is enforced.


|min_throughput
|string
a|Minimum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no minimum
throughput is enforced.
```

```
|min_throughput_iops
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no minimum throughput is enforced. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|min_throughput_mbps
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no minimum throughput is enforced.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
```

```
qos_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|adaptive
|link:#adaptive[adaptive]
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


|name
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDc6aeaf6f29f0252d80acd7ddf8707a7a]]
= Delete a QoS policy
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/qos/policies/{uuid}`#

*Introduced In:* 9.6

Deletes a QoS policy. All QoS workloads associated with the policy are
removed.

== Related ONTAP commands

* `qos policy-group delete`
* `qos adaptive-policy-group delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|* Introduced in: 9.8


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
```

```
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 8454286
| Modifications to these cluster-scoped preset policies is prohibited.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID07ba04ed89aa2f9df335e94995758eb1]]
= Retrieve a QoS policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qos/policies/{uuid}`#

*Introduced In:* 9.6

Retrieves a specific QoS policy.

== Related ONTAP commands
```

```
* `qos policy-group show`
* `qos adaptive-policy-group show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|* Introduced in: 9.8


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|adaptive
|link:#adaptive[adaptive]
```

```
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


|name
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "string",
    "expected_iops_allocation": "string",
    "peak_iops_allocation": "string"
  },
  "fixed": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ]
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```

```
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#adaptive]
[.api-collapsible-fifth-title]
adaptive

Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|absolute_min_iops
|integer
a|Specifies the absolute minimum IOPS that is used as an override when the
expected_iops is less than this value. These floors are not guaranteed on
non-AFF platforms or when FabricPool tiering policies are set.


|block_size
|string
a|Specifies the block size


|expected_iops
|integer
a|Expected IOPS. Specifies the minimum expected IOPS per TB allocated
based on the storage object allocated size. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|expected_iops_allocation
|string
a|Specifies the size to be used to calculate expected IOPS per TB. The
size options are either the storage object allocated space or the storage
object used space.


|peak_iops
|integer
a|Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on
the storage object allocated size or the storage object used size.


|peak_iops_allocation
|string
a|Specifies the size to be used to calculate peak IOPS per TB. The size
options are either the storage object allocated space or the storage
object used space.
```

```
|===


[#fixed]
[.api-collapsible-fifth-title]
fixed

QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|capacity_shared
|boolean
a|Specifies whether the capacities are shared across all objects that use
this QoS policy-group. Default is false.


|max_throughput
|string
a|Maximum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no maximum
throughput is enforced.


|max_throughput_iops
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no maximum throughput is enforced.


|max_throughput_mbps
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no maximum throughput is enforced.


|min_throughput
|string
a|Minimum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no minimum
```

```
throughput is enforced.


|min_throughput_iops
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no minimum throughput is enforced. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|min_throughput_mbps
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no minimum throughput is enforced.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID936cf71c04cda2c85c1b505a73a2c71c]]
= Update a QoS policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/qos/policies/{uuid}`#

*Introduced In:* 9.6

Update a specific QoS policy.

== Related ONTAP commands

* `qos policy-group modify`
* `qos adaptive-policy-group modify`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
```

```
|path
|True
a|* Introduced in: 9.8


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|adaptive
|link:#adaptive[adaptive]
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


|name
```

```
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "adaptive": {
    "block_size": "string",
    "expected_iops_allocation": "string",
    "peak_iops_allocation": "string"
  },
```

```
  "fixed": {
    "capacity_shared": true,
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ]
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "string",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```


Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 8454147
| The maximum limit for QoS policies has been reached.

| 8454154
| The name specified for creating conflicts with an existing QoS policy
name.
```

```
| 8454194
| The minimum throughput value for the policy group must be less than or
equal to the maximum throughput value.

| 8454260
| Invalid value for maximum and minimum fields. Valid values for
max_throughput_iops and max_throughput_mbps combination is for the ratio
of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.

| 8454273
| Invalid value for an adaptive field. Value should be non-zero.

| 8454277
| The name specified for creating an adaptive QoS policy conflicts with an
existing fixed QoS policy name.

| 8454278
| The name specified for creating a fixed QoS policy conflicts with an
existing adaptive QoS policy name.

| 8454286
| Modifications on these cluster scoped preset policies is prohibited.

| 8454327
| The existing fixed QoS policy cannot be modified to an adaptive QoS
policy.

| 8454328
| The existing adaptive QoS policy cannot be modified to a fixed QoS
policy.

| 8454379
| The name specified for creating a fixed QoS policy already exists.

| 8454380
| The name specified for creating an adaptive QoS policy already exists.

| 8454389
| One specified field cannot be used in combination with at least one
other specified field.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#adaptive]
[.api-collapsible-fifth-title]
adaptive

Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|absolute_min_iops
|integer
a|Specifies the absolute minimum IOPS that is used as an override when the
expected_iops is less than this value. These floors are not guaranteed on
non-AFF platforms or when FabricPool tiering policies are set.
```

```
|block_size
|string
a|Specifies the block size


|expected_iops
|integer
a|Expected IOPS. Specifies the minimum expected IOPS per TB allocated
based on the storage object allocated size. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|expected_iops_allocation
|string
a|Specifies the size to be used to calculate expected IOPS per TB. The
size options are either the storage object allocated space or the storage
object used space.


|peak_iops
|integer
a|Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on
the storage object allocated size or the storage object used size.


|peak_iops_allocation
|string
a|Specifies the size to be used to calculate peak IOPS per TB. The size
options are either the storage object allocated space or the storage
object used space.


|===


[#fixed]
[.api-collapsible-fifth-title]
fixed

QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|max_throughput
|string
a|Maximum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no maximum
throughput is enforced.


|max_throughput_iops
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no maximum throughput is enforced.


|max_throughput_mbps
|integer
a|Maximum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no maximum throughput is enforced.


|min_throughput
|string
a|Minimum throughput defined by this policy. It is specified in terms of
Kbps, Mbps or Gbps along with or without IOPS. 0 means no minimum
throughput is enforced.


|min_throughput_iops
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
IOPS. 0 means no minimum throughput is enforced. These floors are not
guaranteed on non-AFF platforms or when FabricPool tiering policies are
set.


|min_throughput_mbps
|integer
a|Minimum throughput defined by this policy. It is specified in terms of
Mbps. 0 means no minimum throughput is enforced.


|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|adaptive
|link:#adaptive[adaptive]
a|Adaptive QoS policy-groups define measurable service level objectives
(SLOs) that adjust based on the storage object used space and the storage
object allocated space.


|fixed
|link:#fixed[fixed]
```

```
a|QoS policy-groups define a fixed service level objective (SLO) for a
storage object.


|name
|string
a|Name of the QoS policy.


|object_count
|integer
a|Number of objects attached to this policy.


|pgid
|integer
a|Policy group ID of the QoS policy.


|policy_class
|string
a|Class of the QoS policy.


|scope
|string
a|Scope of the entity. Set to "cluster" for cluster owned objects and to
"svm" for SVM owned objects.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage QoS options

:leveloffset: +1


[[IDef0bb4c6ca2b5639a8e765e0c6ef1c8c]]
= Manage QoS options



== Quality of Service Options

A QoS option represents a configuration detail that is used by QoS.
 +

'''
```

```
== Examples

=== Retrieving a QoS option from the cluster

The following example retrieves the QoS option in the cluster.
 +

'''

----
curl -X GET "https://<mgmt-ip>/api/storage/qos/qos-
options?return_timeout=0" -H "accept: application/json" -H "Content-Type:
application/json"
----

'''

=== 2) Update a QoS option

The following example shows how to modify the background task reserve
policy to 40%.
 +

'''

----
curl -X PATCH "https://<mgmt-ip>/api/storage/qos/qos-
options?return_timeout=0" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"background_task_reserve\": 40 }"
----

-

'''





[[ID4ad976993ab49eee00ddeb50860e6582]]
= Retrieve QoS options

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qos/qos-options`#

*Introduced In:* 9.14
```

```
Retrieves QoS options.
This option is available only at diagnostic privilege level and above.

== Related ONTAP commands

* `qos settings cluster-options show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|background_task_reserve
|integer
|query
|False
a|Filter by background_task_reserve


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|background_task_reserve
|integer
a|Percentage reserve for critical background tasks.


|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "background_task_reserve": 33
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|
```

```
|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6b2ecf26686d5faaacef5f0285dcbbaa]]
= Update a QoS option

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/qos/qos-options`#

*Introduced In:* 9.14

Update a specific QoS option.
This option is available only at diagnostic privilege level and above.

== Related ONTAP commands

* `qos settings cluster-options modify`
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|background_task_reserve
|integer
a|Percentage reserve for critical background tasks.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "background_task_reserve": 33
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 8454377
| This command is not supported on this platform.

| 8454378
| Unsupported value specified for background-task-reserve. Supported
```

```
values are 0 to 99.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#qos_option]
[.api-collapsible-fifth-title]
qos_option

[cols=3*,options=header]
|===
|Name
|Type
|Description

|background_task_reserve
|integer
```

```
a|Percentage reserve for critical background tasks.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage QoS workloads

:leveloffset: +1


[[IDe5dda138a62eba921260a17d333e6bc2]]
= Manage QoS workloads



== Quality of Service Workloads

A QoS workload represents a storage object that is tracked by QoS.
 +

'''

== Examples

=== Retrieving a list of QoS workloads from the cluster

The following example retrieves all the workloads in the cluster.
 +
```

```
'''

----
curl -X GET "https://<mgmt-ip>/api/storage/qos/workloads" -H "accept:
application/hal+jsON"
----

'''

=== Retrieving a specific QoS workload from the cluster

The following example retrieves a requested workload from the cluster.
 +

'''

----
curl -X GET "https://<mgmt-ip>/api/storage/qos/workloads/77b68b1c-a458-
11eb-baaa-005056bb873e" -H "accept: application/hal+json"
----

'''




[[ID1655b89fd35ec4a7cdfb4cbab1d99cea]]
= Retrieve QoS workloads

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qos/workloads`#

*Introduced In:* 9.10

Retrieves a collection of QoS workloads.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
```

```
|Description

|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|name
|string
|query
|False
a|Filter by name


|wid
|integer
|query
|False
a|Filter by wid


|policy.name
|string
|query
|False
a|Filter by policy.name


|policy.uuid
|string
|query
|False
a|Filter by policy.uuid


|lun
|string
|query
```

```
|False
a|Filter by lun


|volume
|string
|query
|False
a|Filter by volume


|qtree
|string
|query
|False
a|Filter by qtree


|workload_class
|string
|query
|False
a|Filter by workload_class


|uuid
|string
|query
|False
a|Filter by uuid


|file
|string
|query
|False
a|Filter by file


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
```

```
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#qos_workload[qos_workload]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
```

```
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "file": "string",
      "lun": "string",
      "name": "volume1-wid123",
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "qtree": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volume": "volume1",
      "wid": 123,
      "workload_class": "autovolume"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy

QoS policy group reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#qos_workload]
[.api-collapsible-fifth-title]
qos_workload

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|file
|string
a|Name of the file.


|lun
|string
a|Name of the LUN. The name of the LUN will be displayed as "(unknown)" if
the name cannot be retrieved.


|name
|string
a|Name of the QoS workload.
```

```
|policy
|link:#policy[policy]
a|QoS policy group reference.


|qtree
|string
a|Name of the Qtree.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|volume
|string
a|Name of the volume. The name of the volume will be displayed as
"(unknown)" if the name cannot be retrieved.


|wid
|integer
a|Workload ID of the QoS workload.


|workload_class
|string
a|Class of the QoS workload.


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID91317073bef330807117e1b1131d2652]]
= Retrieve a QoS workload

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qos/workloads/{uuid}`#

*Introduced In:* 9.10

Retrieves a specific QoS workload.

== Related ONTAP command

* `qos workload show`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|file
|string
a|Name of the file.


|lun
|string
a|Name of the LUN. The name of the LUN will be displayed as "(unknown)" if
the name cannot be retrieved.
```

```
|name
|string
a|Name of the QoS workload.


|policy
|link:#policy[policy]
a|QoS policy group reference.


|qtree
|string
a|Name of the Qtree.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|volume
|string
a|Name of the volume. The name of the volume will be displayed as
"(unknown)" if the name cannot be retrieved.


|wid
|integer
a|Workload ID of the QoS workload.


|workload_class
|string
a|Class of the QoS workload.


|===


.Example response
[%collapsible%closed]
```

```
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "file": "string",
  "lun": "string",
  "name": "volume1-wid123",
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qtree": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": "volume1",
  "wid": 123,
  "workload_class": "autovolume"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#policy]
[.api-collapsible-fifth-title]
policy

QoS policy group reference.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
```

```
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage qtrees

:leveloffset: +1


[[ID82dd02e41b64de8d8e4203b4a4e0ada8]]
= Manage storage qtrees



== Overview

A qtree is a logically defined file system that can exist as a special
subdirectory of the root directory within a FlexVol volume or a FlexGroup
volume.

== Qtree QoS policy

Qtree QoS policy and settings enforce Service Level Objectives (SLOs) on a
qtree. SLOs can be set by specifying "qos_policy.max_throughput_iops"
and/or "qos_policy.max_throughput_mbps" or
"qos_policy.min_throughput_iops" and/or "qos_policy.min_throughput_mbps".
Specifying "min_throughput_iops" or "min_throughput_mbps" is only
supported on volumes hosted on a node that is flash optimized. A pre-
created QoS policy can also be used by specifying "qos_policy.name" or
"qos_policy.uuid" properties.
Setting or assigning a QoS policy to a qtree is not supported if its
containing volume or SVM has a QoS policy attached, or a file or LUN in
its containing volume already has a QoS policy attached.

== Performance monitoring

Performance of a qtree can be monitored by observing the `statistics.+*+`
properties. These properties show the performance of the qtree in terms of
IOPS and throughput. The `statistics.+*+` properties denote a real-time
monotonically increasing value aggregated across all nodes.

=== Extended performance monitoring
```

```
Extended performance monitoring enables the collection of latency
statistics and the archival of statistics samples for a qtree. When
extended performance monitoring is enabled for a qtree, its performance
can be monitored by observing the `metric.+*+` and `statistics.+*+`
properties. These properties show the performance of the qtree in terms of
IOPS, latency and throughput. The `metric.+*+` properties denote an
average whereas `statistics.+*+` properties denote a real-time
monotonically increasing value aggregated across all nodes.
Extended performance monitoring can be enabled for a qtree using the
`ext_performance_monitoring.enabled` property. A maximum of 50,000 qtrees
can have extended performance monitoring enabled on a cluster. If extended
performance monitoring is enabled for an existing qtree, the prior
`statistics.+*+` properties showing the IOPS and throughput performance of
the qtree will be cleared.

When extended performance monitoring is enabled for a qtree for the first
time in an SVM, existing NFS clients might not be accounted for until
their NFS shares associated with the SVM are remounted.

== Qtree APIs

The following APIs are used to create, retrieve, modify, and delete
qtrees.

&ndash; POST      /api/storage/qtrees

&ndash; GET       /api/storage/qtrees

&ndash; GET       /api/storage/qtrees/\{volume-uuid}/\{qtree-id}

&ndash; PATCH     /api/storage/qtrees/\{volume-uuid}/\{qtree-id}

&ndash; DELETE    /api/storage/qtrees/\{volume-uuid}/\{qtree-id}

== Examples

=== Creating a qtree inside a volume for an SVM

This API is used to create a qtree inside a volume for an SVM.

The following example shows how to create a qtree in a FlexVol volume with
a given security style, user, group, UNIX permissions, an export policy,
and a QoS policy.

'''
```

```
----

# The API:
POST /api/storage/qtrees

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/qtrees?return_records=true' -H
'accept: application/hal+json' -d @test_qtree_post.txt
test_qtree_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "fv"
},
"name": "qt1",
"security_style": "unix",
"user": {
  "name": "unix_user1"
},
"group": {
  "name": "unix_group1"
},
"unix_permissions": 744,
"export_policy": {
    "name": "default"
},
"qos_policy": {
    "max_throughput_iops": 1000
}
}

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "name": "qt1",
      "security_style": "unix",
```

```
      "user": {
        "name": "unix_user1"
      },
      "group": {
        "name": "unix_group1"
      },
      "unix_permissions": 744,
      "export_policy": {
        "name": "default"
      },
      "qos_policy": {
        "min_throughput_iops": 0,
        "min_throughput_mbps": 0,
        "max_throughput_iops": 1000,
        "max_throughput_mbps": 0,
        "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
        "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
      },
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/?volume.name=fv&name=qt1"
        }
      }
    }
  ],
  "job": {
    "uuid": "84edef3c-4f6d-11e9-9a71-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/84edef3c-4f6d-11e9-9a71-005056a7f717"
      }
    }
  }
}
----

'''

=== Retrieving qtrees

This API is used to retrieve qtrees.

The following example shows how to retrieve qtrees belonging to SVM _svm1_
and volume _fv_. The `svm.name` and `volume.name` query parameters are
used to find the required qtrees.
```

```
'''

----

# The API:
GET /api/storage/qtrees

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svm1&volume.name=fv" -H 'accept:
application/hal+json'

# The response
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "id": 0,
      "name": "",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
```

```
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "id": 1,
      "name": "qt1",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/1"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
```

```
      },
      "id": 2,
      "name": "qt2",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/2"
        }
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/?svm.name=svm1&volume.name=fv"
    }
  }
}
----

'''

=== Retrieving properties of a specific qtree using a qtree identifier

This API is used to retrieve properties of a specific qtree using
qtree.id.

The following example shows how to use the qtree identifier to retrieve
all properties of the qtree using the `fields` query parameter.

'''

----

# The API:
GET /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/2?fields=*' -H 'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
```

```
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  },
  "id": 2,
  "name": "qt2",
  "security_style": "unix",
  "user": {
    "name": "unix_user1"
  },
  "group": {
    "name": "unix_group1"
  },
  "unix_permissions": 744,
  "export_policy": {
    "name": "default",
    "id": 12884901889,
    "_links": {
      "self": {
        "href": "/api/protocols/nfs/export-policies/12884901889"
      }
    }
  },
  "qos_policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
    "_links": {
      "self": {
        "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-
005056a7ab52"
      }
    }
```

```
  },
  "ext_performance_monitoring": {
    "enabled": false
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 3,
      "total": 3
    },
    "throughput_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "path": "/fv/qt2",
  "nas": {
    "path": "/fv/qt2",
  },
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
    }
  }
}
----

'''

=== Retrieving properties of a specific qtree using the qtree name

This API is used to retrieve properties of a specific qtree using
"qtree.name".
The following example shows how to retrieve all of the properties
belonging to qtree "qt2". The `svm.name` and `volume.name` query
parameters are used here along with the qtree name.

'''

----
```

```
# The API:
GET /api/storage/qtrees/

# The call:
curl -X GET 'https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svm1&volume.name=fv&name=qt2&fields=*' -H
'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  },
  "id": 2,
  "name": "qt2",
  "security_style": "unix",
  "user": {
    "name": "unix_user1"
  },
  "group": {
    "name": "unix_group1"
  },
  "unix_permissions": 744,
  "export_policy": {
    "name": "default",
    "id": 12884901889,
    "_links": {
      "self": {
        "href": "/api/protocols/nfs/export-policies/12884901889"
      }
    }
  },
```

```
  "qos_policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
    "_links": {
      "self": {
        "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-
005056a7ab52"
      }
    }
  },
  "ext_performance_monitoring": {
    "enabled": false
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 3,
      "total": 3
    },
    "throughput_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
    }
  }
}
----

'''

=== Updating a qtree

This API is used to update a qtree.
```

```
The following example shows how to update properties in a qtree.

'''

----

# The API:
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-
9a71-005056a7f717/2' -H 'accept: application/hal+json' -d
'@test_qtree_patch.txt'
test_qtree_patch.txt(body):
{
"security_style": "mixed",
"user": {
  "name": "unix_user1"
},
"group": {
  "name": "unix_group1"
},
"unix_permissions": 777,
"export_policy": {
    "id": "9",
    "name": "exp1"
},
"qos_policy": {
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab53"
}
}
----

'''

=== Renaming a qtree

This API is used to rename a qtree.

The following example below shows how to rename a qtree with a new name.

'''

----

# The API:
```

```
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-
9a71-005056a7f717/1' -H 'accept: application/hal+json' -d '{ "name":
"new_qt1" }'
----

'''

=== Deleting a qtree inside a volume of an SVM

This API is used to delete a qtree inside a volume of an SVM.

The following example shows how to delete a qtree.

'''

----

# The API:
DELETE /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-
9a71-005056a7f717/2" -H 'accept: application/hal+json'
----

'''

=== Enabling extended performance monitoring on a qtree using PATCH

This API is used to enable extended performance monitoring on a qtree.

The following example shows how to enable extended performance monitoring
on a qtree.

'''

----

# The API:
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-
9a71-005056a7f717/1" -H 'accept: application/hal+json' -d '{
```

```
"ext_performance_monitoring": { "enabled": "true" }}'
----

'''

=== Disabling extended performance monitoring on a qtree using PATCH

This API is used to disable extended performance monitoring on a qtree.

The following example shows how to disable extended performance monitoring
on a qtree.

'''

----

# The API:
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-
9a71-005056a7f717/1" -H 'accept: application/hal+json' -d '{
"ext_performance_monitoring": { "enabled": "false" }}'
----

'''

=== Retrieving performance properties for a qtree which has extended
performance monitoring enabled

The following example shows how to use a GET request to retrieve
performance properties for a qtree which has extended performance
monitoring enabled.

'''

----

# The API:
GET /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/3?fields=ext_performance_monitoring,statistics,metric' -H
'accept: application/hal+json'
{
  "volume": {
```

```
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  },
  "id": 3,
  "name": "qt3",
  "ext_performance_monitoring": {
    "enabled": true
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "throughput_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
    "latency_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "metric": {
    "timestamp": "2019-04-09T05:50:00Z",
    "duration": "PT5M",
    "status": "ok",
    "iops": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
```

```
    "throughput": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "latency": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/3"
    }
  }
}
----

'''
```


```
[[ID36e954c993c0af6de99256a81f1a76ec]]
= Retrieve qtrees

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qtrees`#

*Introduced In:* 9.6

Retrieves qtrees configured for all FlexVol volumes or FlexGroup volumes.

Use the `fields` query parameter to retrieve all properties of the qtree.
If the `fields` query parameter is not used, then GET returns the qtree
`name` and qtree `id` only.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
```

```
g_specific_fields[Requesting specific fields] to learn more.

* `ext_performance_monitoring.enabled`
* `statistics.+*+`
* `metric.+*+`

== Related ONTAP commands

* `qtree show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|id
|integer
|query
|False
a|Filter by id

* Max value: 4994
* Min value: 0


|_tags
```

```
|string
|query
|False
a|Filter by _tags

* Introduced in: 9.13


|security_style
|string
|query
|False
a|Filter by security_style


|name
|string
|query
|False
a|Filter by name


|group.id
|string
|query
|False
a|Filter by group.id

* Introduced in: 9.9


|group.name
|string
|query
|False
a|Filter by group.name

* Introduced in: 9.9


|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other

* Introduced in: 9.16
```

```
|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total

* Introduced in: 9.16


|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write

* Introduced in: 9.16


|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read

* Introduced in: 9.16


|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

* Introduced in: 9.8


|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total

* Introduced in: 9.8
```

```
|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

* Introduced in: 9.8


|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read

* Introduced in: 9.8


|statistics.throughput_raw.other
|integer
|query
|False
a|Filter by statistics.throughput_raw.other

* Introduced in: 9.8


|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total

* Introduced in: 9.8


|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.8


|statistics.throughput_raw.read
|integer
```

```
|query
|False
a|Filter by statistics.throughput_raw.read

* Introduced in: 9.8


|statistics.status
|string
|query
|False
a|Filter by statistics.status

* Introduced in: 9.8


|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp

* Introduced in: 9.8


|ext_performance_monitoring.enabled
|boolean
|query
|False
a|Filter by ext_performance_monitoring.enabled

* Introduced in: 9.16


|path
|string
|query
|False
a|Filter by path


|unix_permissions
|integer
|query
|False
a|Filter by unix_permissions
```

```
|nas.path
|string
|query
|False
a|Filter by nas.path

* Introduced in: 9.9


|volume.name
|string
|query
|False
a|Filter by volume.name


|volume.uuid
|string
|query
|False
a|Filter by volume.uuid


|user.id
|string
|query
|False
a|Filter by user.id

* Introduced in: 9.9


|user.name
|string
|query
|False
a|Filter by user.name

* Introduced in: 9.9


|metric.throughput.other
|integer
|query
|False
a|Filter by metric.throughput.other
```

```
* Introduced in: 9.16


|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total

* Introduced in: 9.16


|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write

* Introduced in: 9.16


|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read

* Introduced in: 9.16


|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp

* Introduced in: 9.16


|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other

* Introduced in: 9.16
```

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total

* Introduced in: 9.16


|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write

* Introduced in: 9.16


|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read

* Introduced in: 9.16


|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other

* Introduced in: 9.16


|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total

* Introduced in: 9.16
```

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write

* Introduced in: 9.16


|metric.latency.read
|integer
|query
|False
a|Filter by metric.latency.read

* Introduced in: 9.16


|metric.status
|string
|query
|False
a|Filter by metric.status

* Introduced in: 9.16


|metric.duration
|string
|query
|False
a|Filter by metric.duration

* Introduced in: 9.16


|qos_policy.min_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.min_throughput_mbps

* Introduced in: 9.8
* Max value: 4194303
* Min value: 0
```

```
|qos_policy.max_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.max_throughput_iops

* Introduced in: 9.8
* Max value: 2147483647
* Min value: 0


|qos_policy.max_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.max_throughput_mbps

* Introduced in: 9.8
* Max value: 4194303
* Min value: 0


|qos_policy.name
|string
|query
|False
a|Filter by qos_policy.name

* Introduced in: 9.8


|qos_policy.min_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.min_throughput_iops

* Introduced in: 9.8
* Max value: 2147483647
* Min value: 0


|qos_policy.uuid
|string
|query
|False
a|Filter by qos_policy.uuid
```

```
* Introduced in: 9.8


|qos_policy.max_throughput
|string
|query
|False
a|Filter by qos_policy.max_throughput

* Introduced in: 9.17


|qos_policy.min_throughput
|string
|query
|False
a|Filter by qos_policy.min_throughput

* Introduced in: 9.17


|export_policy.name
|string
|query
|False
a|Filter by export_policy.name


|export_policy.id
|integer
|query
|False
a|Filter by export_policy.id


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
```

```
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Max value: 120
* Min value: 0
* Default value: 15


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#qtree[qtree]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "id": 100,
        "name": "default"
      },
      "group": {
        "id": "20001",
        "name": "unix_group1"
      },
      "id": 1,
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT5M",
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "name": "string",
      "nas": {
        "path": "/volume3/qtree1"
      },
      "path": "/volume3/qtree1",
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput": [
```

```
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "security_style": "string",
      "statistics": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "svm": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_permissions": 493,
      "user": {
        "id": "10001",
        "name": "unix_user1"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 918235
| A volume with UUID was not found.

| 2621462
| The specified SVM does not exist.

| 5242889
| Failed to get the qtree from volume.

| 5242956
| Failed to obtain qtree.
```

```
| 5242965
| Invalid qtree path. The volume name component of the qtree path, must be
the same as the volume specified with the parameter.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#export_policy]
[.api-collapsible-fifth-title]
export_policy

Export Policy


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|

|name
|string
a|

|===


[#ext_performance_monitoring]
[.api-collapsible-fifth-title]
ext_performance_monitoring

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|enabled
|boolean
a|Specifies whether extended performance monitoring is enabled for the
qtree.


|===


[#group]
[.api-collapsible-fifth-title]
group

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the group that owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric group name of group that owns the qtree. Valid in POST or
PATCH.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
```

```
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

```
When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|max_throughput
|string
a|Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none. This is
mutually exclusive with name and UUID during POST and PATCH. This cannot
be set when either max_throughput_mbps or max_throughput_iops are set
during POST or PATCH. During GET, the returned value is rounded to the
largest unit with a value greater than 1.


|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually
exclusive with name and UUID during POST and PATCH. This cannot be set
when max_throughput is set during POST or PATCH.


|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH. This
cannot be set when max_throughput is set during POST or PATCH.


|min_throughput
|string
```

```
a|Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none.Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when either min_throughput_mbps
or min_throughput_iops are set during POST or PATCH. During GET, the
returned value is rounded to the largest unit with a value greater than 1.


|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when min_throughput is set
during POST or PATCH.


|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.This
cannot be set when min_throughput is set during POST or PATCH.


|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency observed at the storage object, in microseconds. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
```

```
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency observed at the storage object, in microseconds. This
can be divided by the raw IOPS value to calculate the average latency per
I/O operation.


|status
|string
```

```
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled with the next closest collection and tagged with
"backfilled_data". "inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "inconsistent_old_data" is returned when one or more nodes does not
have the latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

```
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#user]
[.api-collapsible-fifth-title]
user

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the user who owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric username of user who owns the qtree. Valid in POST or
PATCH.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

Required in POST


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

A qtree is a directory at the top level of a volume to which a custom
export policy (for fine-grained access control) and a quota rule can be
applied, if required.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.


|nas
|link:#nas[nas]
a|

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.
```

```
|qos_policy
|link:#qos_policy[qos_policy]
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|statistics
|link:#statistics[statistics]
a|These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.


|svm
|link:#svm[svm]
a|Required in POST


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|volume
|link:#volume[volume]
a|Required in POST


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID1a751aae7a133dd3630a2275b0f477e3]]
= Create a qtree in a FlexVol or FlexGroup volume

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/qtrees`#

*Introduced In:* 9.6

Creates a qtree in a FlexVol volume or a FlexGroup volume.

After a qtree is created, the new qtree is assigned an identifier. This
identifier is obtained using a qtree GET request. This identifier is used
in the API path for the qtree PATCH and DELETE operations.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the qtree.
* `volume.uuid` or `volume.name` - Existing volume in which to create the
qtree.
* `name` - Name for the qtree.

== Recommended optional properties

If not specified in POST, the values are inherited from the volume.

* `security_style` - Security style for the qtree.
* `unix_permissions` - UNIX permissions for the qtree.
* `export_policy.name or export_policy.id` - Export policy of the SVM for
the qtree.

== Related ONTAP commands

* `qtree create`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.


|nas
|link:#nas[nas]
a|

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.


|qos_policy
|link:#qos_policy[qos_policy]
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
```

```
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|svm
|link:#svm[svm]
a|Required in POST


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|volume
|link:#volume[volume]
a|Required in POST


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
```

```
    "id": 100,
    "name": "default"
  },
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "name": "string",
  "nas": {
    "path": "/volume3/qtree1"
  },
  "path": "/volume3/qtree1",
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "user": {
```

```
    "id": "10001",
    "name": "unix_user1"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 262245
| Invalid field value.

| 262247
| Invalid field value.

| 917525
| The specified volume does not exist in Vserver.

| 917927
| The specified volume was not found.

| 918232
| Either `volume.name` or `volume.uuid` must be provided.

| 918236
| The specified `volume.uuid` and `volume.name` refer to different
volumes.

| 1703954
| Export Policy name specified is invalid.

| 2621462
| The specified SVM does not exist.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be provided.

| 5242881
| Cannot create qtree because the volume is read-only.

| 5242886
| Failed to create qtree.
```

```
| 5242894
| Qtree with empty name "" is not allowed, as that is reserved for the
default qtree.

| 5242900
| Qtree not supported on FlexCache volume

| 5242948
| Qtree is not supported on FlexCache origin volume.

| 5242951
| Export Policy supplied does not belong to the specified Export Policy
ID.

| 5242952
| Export Policy ID specified is invalid.

| 5242953
| Qtree name must be provided.

| 5242967
| UNIX user or group ID must be 32-bit unsigned integer.

| 5242970
| FlexCache create is in progress for the volume.

| 5242978
| The maximum number of qtrees for which extended performance monitoring
can be enabled has been reached. Retry the POST request with
ext_performance_monitoring.enabled set to false.

| 5242979
| Qtree creation succeeded but failed to enable extended performance
monitoring on the qtree.

| 6622064
| Security-style NTFS is not supported on a SnapMirror active sync
relationship volume.

| 8454348
| QoS on qtrees is not supported because not all nodes in the cluster can
support it.

| 9437324
| The security style unified is not supported.
```

```
| 23724050
| Failed to resolve user or group name.

| 66846755
| Failed to determine whether volume is a FlexCache volume or not.

| 66846839
| Failed to determine the effective cluster version of all the nodes
hosting FlexCache volumes connected to FlexCache origin volume.

| 92405926
| Qtree operation failed on the specified object store volume.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#export_policy]
```

```
[.api-collapsible-fifth-title]
export_policy

Export Policy


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|name
|string
a|

|===


[#ext_performance_monitoring]
[.api-collapsible-fifth-title]
ext_performance_monitoring

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether extended performance monitoring is enabled for the
qtree.


|===


[#group]
[.api-collapsible-fifth-title]
group

The user set as owner of the qtree.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the group that owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric group name of group that owns the qtree. Valid in POST or
PATCH.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
```

```
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.
```

```
|===


[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_throughput
|string
a|Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none. This is
```

```
mutually exclusive with name and UUID during POST and PATCH. This cannot
be set when either max_throughput_mbps or max_throughput_iops are set
during POST or PATCH. During GET, the returned value is rounded to the
largest unit with a value greater than 1.


|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually
exclusive with name and UUID during POST and PATCH. This cannot be set
when max_throughput is set during POST or PATCH.


|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH. This
cannot be set when max_throughput is set during POST or PATCH.


|min_throughput
|string
a|Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none.Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when either min_throughput_mbps
or min_throughput_iops are set during POST or PATCH. During GET, the
returned value is rounded to the largest unit with a value greater than 1.


|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when min_throughput is set
during POST or PATCH.


|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.This
cannot be set when min_throughput is set during POST or PATCH.
```

```
|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
```

```
|integer
a|Performance metric for write I/O operations.


|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency observed at the storage object, in microseconds. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency observed at the storage object, in microseconds. This
can be divided by the raw IOPS value to calculate the average latency per
I/O operation.


|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled with the next closest collection and tagged with
"backfilled_data". "inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "inconsistent_old_data" is returned when one or more nodes does not
have the latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.


|timestamp
|string
a|The timestamp of the performance data.
```

```
|===


[#svm]
[.api-collapsible-fifth-title]
svm

Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#user]
[.api-collapsible-fifth-title]
user

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the user who owns the qtree. Valid in POST or PATCH.
```

```
|name
|string
a|Alphanumeric username of user who owns the qtree. Valid in POST or
PATCH.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree
```

```
A qtree is a directory at the top level of a volume to which a custom
export policy (for fine-grained access control) and a quota rule can be
applied, if required.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.


|nas
|link:#nas[nas]
a|

|path
|string
```

```
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.


|qos_policy
|link:#qos_policy[qos_policy]
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|svm
|link:#svm[svm]
a|Required in POST


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|volume
|link:#volume[volume]
a|Required in POST


|===


[#job_link]
[.api-collapsible-fifth-title]
```

```
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID851234555765ae8468add38ebf442dbf]]
= Delete a qtree

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/qtrees/{volume.uuid}/{id}`#

*Introduced In:* 9.6

Deletes a qtree.

== Related ONTAP commands

* `qtree delete`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|id
|string
|path
|True
a|Qtree ID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 917505
| Vserver not found.
```

```
| 917506
| Volume not found.

| 917525
| The specified volume does not exist in Vserver.

| 918235
| A volume with UUID was not found.

| 5242894
| The default qtree cannot be deleted.

| 5242895
| Failed to delete the qtree.

| 5242897
| This operation is not permitted on read-only volume.

| 5242898
| This operation is only permitted on a data Vserver.

| 5242916
| Cannot delete qtree because the volume contains one or more LUNs.

| 5242925
| The limit for the number of concurrent delete jobs has been reached.

| 5242927
| Unable to find qtree.

| 5242955
| The UUID of the volume is required.

| 5242957
| Failed to delete qtree with ID in the volume and SVM.

| 5242965
| Invalid qtree path. The volume name component of the qtree path, must be
the same as the volume specified with the parameter.

| 10485796
| Cannot delete qtree because it contains a Storage Level Access Guard
(SLAG).
|===

Also see the table of common errors in the
```

```
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
```

```
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb1d6d60f3daaf94a933ddb2c063894a5]]
= Retrieve qtree properties

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qtrees/{volume.uuid}/{id}`#

*Introduced In:* 9.6

Retrieves properties for a specific qtree identified by the `volume.uuid`
and the `id` in the API path.

== Expensive properties

There is an added computational cost to retrieving values for these
properties. They are not included by default in GET results and must be
explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `ext_performance_monitoring.enabled`
* `statistics.+*+`
* `metric.+*+`

== Related ONTAP commands

* `qtree show`
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|id
|string
|path
|True
a|Qtree ID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

```
|link:#_links[_links]
a|

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.


|nas
|link:#nas[nas]
a|

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.
```

```
|qos_policy
|link:#qos_policy[qos_policy]
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|statistics
|link:#statistics[statistics]
a|These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.


|svm
|link:#svm[svm]
a|Required in POST


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|volume
|link:#volume[volume]
a|Required in POST


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT5M",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
```

```
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "string",
  "nas": {
    "path": "/volume3/qtree1"
  },
  "path": "/volume3/qtree1",
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
```

```
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "user": {
    "id": "10001",
    "name": "unix_user1"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 918235
| A volume with UUID was not found.

| 2621462
| The specified SVM does not exist.

| 5242889
| Failed to get the qtree from volume.

| 5242956
| Failed to obtain a qtree with ID.

| 5242965
| Invalid qtree path. The volume name component of the qtree path, must be
the same as the volume specified with the parameter.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#export_policy]
[.api-collapsible-fifth-title]
export_policy

Export Policy


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|

|name
|string
a|

|===


[#ext_performance_monitoring]
[.api-collapsible-fifth-title]
ext_performance_monitoring

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|enabled
|boolean
a|Specifies whether extended performance monitoring is enabled for the
qtree.


|===


[#group]
[.api-collapsible-fifth-title]
group

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the group that owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric group name of group that owns the qtree. Valid in POST or
PATCH.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
```

```
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

```
When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|max_throughput
|string
a|Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none. This is
mutually exclusive with name and UUID during POST and PATCH. This cannot
be set when either max_throughput_mbps or max_throughput_iops are set
during POST or PATCH. During GET, the returned value is rounded to the
largest unit with a value greater than 1.


|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually
exclusive with name and UUID during POST and PATCH. This cannot be set
when max_throughput is set during POST or PATCH.


|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH. This
cannot be set when max_throughput is set during POST or PATCH.


|min_throughput
|string
a|Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec
```

```
or Gigabytes per sec along with or without IOPS. 0 means none.Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when either min_throughput_mbps
or min_throughput_iops are set during POST or PATCH. During GET, the
returned value is rounded to the largest unit with a value greater than 1.


|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when min_throughput is set
during POST or PATCH.


|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.This
cannot be set when min_throughput is set during POST or PATCH.


|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency observed at the storage object, in microseconds. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency observed at the storage object, in microseconds. This
can be divided by the raw IOPS value to calculate the average latency per
I/O operation.


|status
|string
a|Any errors associated with the sample. For example, if the aggregation
```

```
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled with the next closest collection and tagged with
"backfilled_data". "inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "inconsistent_old_data" is returned when one or more nodes does not
have the latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#user]
[.api-collapsible-fifth-title]
user

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the user who owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric username of user who owns the qtree. Valid in POST or
PATCH.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

Required in POST


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID0aeafc3bdde8c18ab0b0721014e1d311]]
= Update properties for a qtree

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/qtrees/{volume.uuid}/{id}`#
```

```
*Introduced In:* 9.6

Updates properties for a specific qtree.

== Related ONTAP commands

* `qtree modify`
* `qtree rename`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|id
|string
|path
|True
a|Qtree ID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.
```

```
* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.
```

```
|nas
|link:#nas[nas]
a|

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.


|qos_policy
|link:#qos_policy[qos_policy]
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_tags": [
```

```
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
    "id": 100,
    "name": "default"
  },
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "name": "string",
  "nas": {
    "path": "/volume3/qtree1"
  },
  "path": "/volume3/qtree1",
  "qos_policy": {
    "max_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput": [
      "900KB/s",
      "500MB/s",
      "120GB/s",
      "5000IOPS",
      "5000IOPS,500KB/s",
      "2500IOPS,100MB/s",
      "1000IOPS,25MB/s"
    ],
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
  "unix_permissions": 493,
  "user": {
```

```
    "id": "10001",
    "name": "unix_user1"
  }
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
| Error Code | Description

| 262196
| The request contains a field which cannot be set in this operation.

| 262278
| Required field is missing in the request.

| 917505
| Vserver not found.

| 917525
| The specified volume does not exist in Vserver.

| 918235
| A volume with UUID was not found.

| 5242887
| Failed to modify qtree.

| 5242897
| This operation is not permitted on read-only volume.

| 5242898
| This operation is only permitted on a data Vserver.

| 5242902
| Missing inputs.

| 5242915
| Failed to assign qtree export policy to qtree.

| 5242927
| Unable to find qtree.

| 5242945
| Failed to modify qtree.

| 5242951
| Export policy supplied does not belong to the specified export policy
ID.

| 5242954
| Failed to get the qtree from volume.

| 5242955
```

```
| The UUID of the volume is required.

| 5242956
| Failed to obtain a qtree with ID.

| 5242957
| Failed to delete the qtree.

| 5242958
| Failed to rename the qtree with ID in the volume and SVM.

| 5242959
| Successfully renamed qtree but the modify operation failed.

| 5242965
| Invalid qtree path. The volume name component of the qtree path, must be
the same as the volume specified with the parameter.

| 5242967
| UNIX user or group ID must be 32-bit unsigned integer.

| 5242971
| Qtree was renamed. However, the path modification failed.

| 5242972
| Cannot rename qtree as that name already exists on a volume in the
Vserver.

| 5242973
| Cannot rename qtree to name with path concurrently on volume in Vserver,
unless non-root qtrees in enabled on the volume.

| 5242974
| Moved qtree. However, other properties were not modified.

| 5242975
| Renamed qtree and moved the qtree. However, other properties were not
modified.

| 6622064
| Security-style NTFS is not supported on a SnapMirror active sync
relationship volume.

| 8454348
| QoS on qtrees is not supported because not all nodes in the cluster can
support it.
```

```
| 9437324
| The security style unified is not supported.

| 23724050
| Failed to resolve user or group name.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#export_policy]
[.api-collapsible-fifth-title]
export_policy

Export Policy


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|id
|integer
a|

|name
|string
a|

|===


[#ext_performance_monitoring]
[.api-collapsible-fifth-title]
ext_performance_monitoring

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether extended performance monitoring is enabled for the
qtree.


|===


[#group]
[.api-collapsible-fifth-title]
group

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|id
|string
a|The numeric ID of the group that owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric group name of group that owns the qtree. Valid in POST or
PATCH.


|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.
```

```
|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
```

```
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.


|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#nas]
[.api-collapsible-fifth-title]
nas
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH.


|===


[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_throughput
|string
a|Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none. This is
mutually exclusive with name and UUID during POST and PATCH. This cannot
be set when either max_throughput_mbps or max_throughput_iops are set
during POST or PATCH. During GET, the returned value is rounded to the
largest unit with a value greater than 1.


|max_throughput_iops
```

```
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually
exclusive with name and UUID during POST and PATCH. This cannot be set
when max_throughput is set during POST or PATCH.


|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH. This
cannot be set when max_throughput is set during POST or PATCH.


|min_throughput
|string
a|Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec
or Gigabytes per sec along with or without IOPS. 0 means none.Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when either min_throughput_mbps
or min_throughput_iops are set during POST or PATCH. During GET, the
returned value is rounded to the largest unit with a value greater than 1.


|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH. This cannot be set when min_throughput is set
during POST or PATCH.


|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.This
cannot be set when min_throughput is set during POST or PATCH.


|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
```

```
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency observed at the storage object, in microseconds. This can
be divided by the raw IOPS value to calculate the average latency per I/O
operation.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#statistics]
[.api-collapsible-fifth-title]
statistics

These are raw IOPS and throughput performance numbers. These numbers are
aggregated across all nodes in the cluster and increase with the uptime of
the cluster.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
```

```
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.


|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency observed at the storage object, in microseconds. This
can be divided by the raw IOPS value to calculate the average latency per
I/O operation.


|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled with the next closest collection and tagged with
"backfilled_data". "inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "inconsistent_old_data" is returned when one or more nodes does not
have the latest data.


|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.


|timestamp
|string
a|The timestamp of the performance data.


|===


[#svm]
[.api-collapsible-fifth-title]
svm
```

```
Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#user]
[.api-collapsible-fifth-title]
user

The user set as owner of the qtree.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The numeric ID of the user who owns the qtree. Valid in POST or PATCH.


|name
|string
a|Alphanumeric username of user who owns the qtree. Valid in POST or
PATCH.
```

```
|===


[#volume]
[.api-collapsible-fifth-title]
volume

Required in POST


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

A qtree is a directory at the top level of a volume to which a custom
export policy (for fine-grained access control) and a quota rule can be
applied, if required.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_tags
|array[string]
a|Tags are an optional way to track the uses of a resource. Tag values
must be formatted as key:value strings.


|export_policy
|link:#export_policy[export_policy]
a|Export Policy


|ext_performance_monitoring
|link:#ext_performance_monitoring[ext_performance_monitoring]
a|

|group
|link:#group[group]
a|The user set as owner of the qtree.


|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.


|name
|string
a|The name of the qtree. Required in POST; optional in PATCH.


|nas
|link:#nas[nas]
a|

|path
|string
a|Client visible path to the qtree. This field is not available if the
volume does not have a junction-path configured. Not valid in POST or
PATCH. This field is to be deprecated and replaced with nas.path.


|qos_policy
|link:#qos_policy[qos_policy]
```

```
a|When "min_throughput_iops", "min_throughput_mbps", "min_throughput",
"max_throughput_iops", "max_throughput_mbps" or "max_throughput"
attributes are specified, the storage object is assigned to an auto-
generated QoS policy group. If the attributes are later modified, the
auto-generated QoS policy-group attributes are modified. Attributes can be
removed by specifying "0" and policy group by specifying "none". Upon
deletion of the storage object or if the attributes are removed, then the
QoS policy-group is also removed.


|security_style
|string
a|Security style. Valid in POST or PATCH.


|unix_permissions
|integer
a|The UNIX permissions for the qtree. Valid in POST or PATCH.


|user
|link:#user[user]
a|The user set as owner of the qtree.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID9e60b08a63d4030cd405cd7bed793d44]]
= Retrieve historical performance metrics for a qtree with extended
performance monitoring enabled

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/qtrees/{volume.uuid}/{qtree.id}/metrics`#

*Introduced In:* 9.16

Retrieves historical performance metrics for a qtree which has extended
performance monitoring enabled.


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|duration
|string
|query
|False
a|Filter by duration


|qtree.name
|string
```

```
|query
|False
a|Filter by qtree.name


|volume.name
|string
|query
|False
a|Filter by volume.name


|status
|string
|query
|False
a|Filter by status


|timestamp
|string
|query
|False
a|Filter by timestamp


|throughput.other
|integer
|query
|False
a|Filter by throughput.other


|throughput.total
|integer
|query
|False
a|Filter by throughput.total


|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read


|iops.other
|integer
|query
|False
a|Filter by iops.other


|iops.total
|integer
|query
|False
a|Filter by iops.total


|iops.write
|integer
|query
|False
a|Filter by iops.write


|iops.read
|integer
|query
|False
a|Filter by iops.read


|latency.other
|integer
|query
|False
a|Filter by latency.other


|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.write
|integer
|query
|False
a|Filter by latency.write


|latency.read
|integer
|query
|False
a|Filter by latency.read


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|volume.uuid
|string
|path
|True
a|Volume UUID


|qtree.id
|string
|path
|True
a|Qtree ID


|interval
|string
|query
```

```
|False
a|The time range for the data. Examples can be 1d, 1m, 1w, 1y.
The period for each time range is as follows:

* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1d", "1w", "1m", "1y"]


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
```

```
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT5M",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 1,
        "name": "qt1"
      },
      "status": "ok",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "read": 200,
        "total": 1000,
```

```
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 918235
| A volume with UUID was not found.

| 2621462
| The specified SVM does not exist.

| 5242956
| Failed to obtain a qtree with ID.

| 5242986
| Extended performance monitoring is not enabled on the qtree.

| 8586227
| Interval not supported for this object.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===


[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
|string
a|The name of the qtree.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.


|read
|integer
a|Performance metric for read I/O operations.


|total
|integer
a|Performance metric aggregated over all types of I/O operations.


|write
|integer
a|Performance metric for write I/O operations.


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

```
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:


|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.


|qtree
|link:#qtree[qtree]
a|

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_ delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.


|timestamp
|string
a|The timestamp of the performance data.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Storage quota

:leveloffset: +1


[[IDea6963e7e80eef5b403759c7d0763081]]
= ONTAP REST API storage quota endpoints



== Overview

Quotas provide a way to restrict or track the files and space usage by a
user, group, or qtree. Quotas are enabled for a specific FlexVol or a
FlexGroup volume.

Quotas can have soft or hard limits. Soft limits cause ONTAP to send a
notification when specified limits are exceeded. Hard limits prevent a
write operation from succeeding when specified limits are exceeded.

== Quota policy rule APIs

Quotas are defined as quota policy rules specific to FlexVol or FlexGroup
volumes. Each quota rule has a type. The type can be either "user",
"group", or "tree".

The following APIs can be used to perform create, retrieve, modify, and
delete operations related to quota policy rules for a FlexVol or a
FlexGroup volume.

&ndash; POST      /api/storage/quota/rules
```

```
&ndash; GET       /api/storage/quota/rules

&ndash; GET       /api/storage/quota/rules/\{rule-uuid}

&ndash; PATCH     /api/storage/quota/rules/\{rule-uuid}

&ndash; DELETE    /api/storage/quota/rules/\{rule-uuid}

== Enabling and disabling quotas

After the quota rules are created, the quota rules need to be enabled on
each FlexVol or FlexGroup volume for soft or hard limits to take effect in
the filesystem. Enabling quotas can be done on a volume-by-volume basis.

The following APIs can be used to enable and disable and obtain the quota
state for a FlexVol or a FlexGroup volume.

&ndash; PATCH  /api/storage/volumes/\{volume-uuid} -d
'{"quota.enabled":"true"}'

&ndash; PATCH  /api/storage/volumes/\{volume-uuid} -d
'{"quota.enabled":"false"}'

&ndash; GET    /api/storage/volumes/\{volume-uuid}/?fields=quota.state

== Quota report APIs

Quota report records provide usage information for a user, group, or qtree
against the quota limits configured on a FlexVol or a FlexGroup volume.
The following APIs can be used to retrieve quota reports associated with a
FlexVol or a FlexGroup volume.

&ndash; GET    /api/storage/quota/reports

&ndash; GET    /api/storage/quota/reports/\{volume-uuid}/\{index}

== Quota resize

Quota resize allows you to modify the quota limits directly in the
filesystem.

It is important to note that quota must be enabled on a FlexVol or a
FlexGroup volume for `quota resize` to take effect.
You can perform a `quota resize` using the quota policy rule PATCH API. If
the quota is disabled on the volume, the quota policy rule PATCH API
modifies the rule, but this does not affect the limits in the filesystem.
```

```
The following API can be used to perform `quota resize` provided quota is
enabled on the volume.

&ndash; PATCH  /api/storage/quota/rules/\{rule-uuid}
The changed limits in the filesystem can be confirmed using the quota
report REST API.

&ndash; GET  /api/storage/quota/reports

== Quota size fields

Quota size fields (e.g. the space.hard_limit property of
/storage/quota/reports) may be set to "`null`" to indicate that there is
no enforced limit on the property. This is equivalent to a value of '-' in
the ONTAP CLI and ONTAPI interfaces.

For querying size fields, setting <field>=null matches properties that
have no limit specified.

For example, to return all records with a files.hard_limit value > 5 OR
files.hard_limit=null, the query command to use is
"<path>?files.hard_limit\=>5\|null".

When ordering the records returned based on size fields, a record whose
property value is null will compare less than a record with a defined
property value. This behavior differs from the ONTAPI API ordering but is
consistent with the handling of integer fields in other REST API
endpoints.





:leveloffset: -1


= Manage storage quota reports

:leveloffset: +1


[[IDe8c4b9e527ae1e98b94af032c7906e36]]
= Manage storage quota reports



== Overview
```

```
Quota reports provide the current file and space consumption for a user,
group, or qtree in a FlexVol or a FlexGroup volume.

== Quota report APIs

The following APIs can be used to retrieve quota reports associated with a
volume in ONTAP.

&ndash; GET       /api/storage/quota/reports

&ndash; GET       /api/storage/quota/reports/\{volume_uuid}/\{index}

== Examples

=== Retrieving all the quota report records

This API is used to retrieve all the quota report records.

The following example shows how to retrieve quota report records for all
FlexVol volumes and FlexGroup volumes.

'''

----

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports"  -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
```

```
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 0,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 1152921504606846976,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/1152921504606846976"
        }
      }
    },
    {
```

```
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 3458764513820540928,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/3458764513820540928"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
```

```
      },
      "index": 4611686018427387904,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/4611686018427387904"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 5764607523034234880,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/5764607523034234880"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
```

```
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 0,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 281474976710656,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/281474976710656"
        }
      }
```

```
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 1152921504606846976,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1152921504606846976"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
```

```
          }
        }
      },
      "index": 1153202979583557632,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1153202979583557632"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 2305843013508661248,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/2305843013508661248"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
```

```
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 3458764513820540928,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/3458764513820540928"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 3459045988797251584,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/3459045988797251584"
```

```
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 4611686018427387904,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/4611686018427387904"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
```

```
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 4611967493404098560,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/4611967493404098560"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 5764607523034234880,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/5764607523034234880"
        }
      }
    }
  ],
  "num_records": 15,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/"
```

```
    }
  }
}
----

'''

=== Retrieving a specific quota report record

This API is used to retrieve a specific quota report record.

The following example shows how to retrieve a single quota report user
record.

'''

----

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-
11e9-8513-005056a7657c/281474976710656"  -H 'accept: application/hal+json'

# Response for quota report user record:
{
"svm": {
  "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
    }
  }
},
"volume": {
  "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
    }
  }
},
"index": 281474976710656,
```

```
"type": "user",
"users": [
  {
    "name": "fred",
    "id"  : "300008"
  }
],
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/281474976710656"
  }
}
}
----

'''

=== Retrieving a single quota report multi-user record
```

```
'''

----

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-
11e9-8513-005056a7657c/281474976710656"  -H 'accept: application/hal+json'

# Response for quota report multi-user record:
{
"svm": {
  "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
    }
  }
},
"volume": {
  "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
    }
  }
},
"index": 1153484454560268288,
"type": "user",
"users": [
  {
    "name": "fred",
    "id"  : "300008"
  },
  {
    "name": "john",
    "id"  : "300009"
  },
  {
    "name": "smith",
    "id"  : "300010"
```

```
    }
  ],
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/1153484454560268288"
    }
  }
}
----

'''

=== Retrieving a single quota report group record

'''

----
```

```
# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-
11e9-8513-005056a7657c/3459045988797251584"  -H 'accept:
application/hal+json'

# Response for quota report group record:
{
"svm": {
  "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
    }
  }
},
"volume": {
  "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
    }
  }
},
"index": 3459045988797251584,
"type": "group",
"group": {
  "name"  : "test_group",
  "id"    : "500009"
},
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
```

```
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/3459045988797251584"
  }
}
}
----

'''

=== Retrieving a single quota report tree record

'''

----

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-
11e9-8513-005056a7657c/4612248968380809216"  -H 'accept:
application/hal+json'

# Response for quota report tree record:
{
"svm": {
  "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
  "name": "svm1",
  "_links": {
```

```
    "self": {
      "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
    }
  }
},
"volume": {
  "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
    }
  }
},
"index": 4612248968380809216,
"type": "tree",
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
```

```
005056a7657c/4612248968380809216"
  }
}
}
----

'''

=== Retrieving only records enforced by non-default rules

'''

----

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?show_default_records=false"  -H 'accept:
application/hal+json'

# Response from only  non-default records
{
"records": [
  {
    "svm": {
      "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
        }
      }
    },
    "volume": {
      "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-
005056a7657c"
        }
      }
    },
    "index": 4612248968380809216,
    "type": "tree",
    "qtree": {
      "name": "qt1",
```

```
      "id": 1,
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-
005056a7657c/1"
        }
      }
    },
    "space": {
      "hard_limit": 41943040,
      "soft_limit": 31457280,
      "used": {
        "total": 10567680,
        "soft_limit_percent": 34,
        "hard_limit_percent": 25
      }
    },
    "files": {
      "soft_limit": 30,
      "hard_limit": 40,
      "used": {
        "total": 11,
        "soft_limit_percent": 37,
        "hard_limit_percent": 28
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/4612248968380809216"
      }
    }
  },
  {
    "svm": {
      "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
        }
      }
    },
    "volume": {
      "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
      "name": "fv",
```

```
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-
005056a7657c"
        }
      }
    },
    "index": 1153484454560268288,
    "type": "user",
    "users": [
      {
        "name": "fred",
        "id"  : "300008"
      },
      {
        "name": "john",
        "id"  : "300009"
      },
      {
        "name": "smith",
        "id"  : "300010"
      }
    ],
    "qtree": {
      "name": "qt1",
      "id": 1,
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-
005056a7657c/1"
        }
      }
    },
    "space": {
      "hard_limit": 41943040,
      "soft_limit": 31457280,
      "used": {
        "total": 10567680,
        "soft_limit_percent": 34,
        "hard_limit_percent": 25
      }
    },
    "files": {
      "soft_limit": 30,
      "hard_limit": 40,
      "used": {
```

```
        "total": 11,
        "soft_limit_percent": 37,
        "hard_limit_percent": 28
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/1153484454560268288"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?show_default_records=false"
  }
}
}
----

'''

=== Retrieving quota report records with query parameters

The following example shows how to retrieve tree type quota report
records.

'''

----

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports?type=tree"  -H
'accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
```

```
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "volume": {
      "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
00505682cd5c"
        }
      }
    },
    "index": 2305843013508661248,
    "type": "tree",
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/2305843013508661248"
      }
    }
  },
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "volume": {
      "uuid": "a5ceebd2-6ccf-11ea-bc35-005056823577",
      "name": "fg",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/a5ceebd2-6ccf-11ea-bc35-
005056823577"
        }
      }
    },
```

```
    "index": 2305843013508661248,
    "type": "tree",
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/a5ceebd2-6ccf-11ea-bc35-
005056823577/2305843013508661248"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?type=tree"
  }
}
}
----

'''

=== Retrieving all the quota reports of a specific volume and the files
fields

'''

----

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?volume.name=fv&fields=files"  -H 'accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
```

```
        }
      }
    },
    "volume": {
      "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
00505682cd5c"
        }
      }
    },
    "index": 410328290557952,
    "files": {
      "soft_limit": 20,
      "hard_limit": 30,
      "used": {
        "total": 0,
        "soft_limit_percent": 0,
        "hard_limit_percent": 0
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/410328290557952"
      }
    }
  },
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "volume": {
      "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
```

```
00505682cd5c"
        }
      }
    },
    "index": 2305843013508661248,
    "files": {
      "soft_limit": 200,
      "hard_limit": 400,
      "used": {
        "total": 4,
        "soft_limit_percent": 2,
        "hard_limit_percent": 1
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/2305843013508661248"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?volume.name=fv&fields=files"
  }
}
}
----

'''

=== Retrieving quota reports for all volumes with the property
files.hard_limit greater than 5 or null (unlimited) for qtree qt1

'''

----

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?qtree.name=qt1&files.hard_limit=>5&#124;null
```

```
" -H 'accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume",
      "uuid": "0c9cff59-da3c-11ed-930f-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0c9cff59-da3c-11ed-930f-
005056ac3135"
        }
      }
    },
    "index": 281474976710656,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/0c9cff59-da3c-11ed-930f-
005056ac3135/281474976710656"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
```

```
    "volume": {
      "name": "srcVolume",
      "uuid": "0c9cff59-da3c-11ed-930f-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0c9cff59-da3c-11ed-930f-
005056ac3135"
        }
      }
    },
    "index": 1153202979583557632,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/0c9cff59-da3c-11ed-930f-
005056ac3135/1153202979583557632"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume",
      "uuid": "0c9cff59-da3c-11ed-930f-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0c9cff59-da3c-11ed-930f-
005056ac3135"
        }
      }
    },
    "index": 2305843013508661248,
    "qtree": {
      "name": "qt1"
    },
    "files": {
```

```
      "hard_limit": 15
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/0c9cff59-da3c-11ed-930f-
005056ac3135/2305843013508661248"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume",
      "uuid": "0c9cff59-da3c-11ed-930f-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0c9cff59-da3c-11ed-930f-
005056ac3135"
        }
      }
    },
    "index": 3459045988797251584,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/0c9cff59-da3c-11ed-930f-
005056ac3135/3459045988797251584"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
```

```
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume",
      "uuid": "0c9cff59-da3c-11ed-930f-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0c9cff59-da3c-11ed-930f-
005056ac3135"
        }
      }
    },
    "index": 4611967493404098560,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/0c9cff59-da3c-11ed-930f-
005056ac3135/4611967493404098560"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume2",
      "uuid": "3a2a8d78-ded3-11ed-9806-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3a2a8d78-ded3-11ed-9806-
005056ac3135"
        }
      }
    },
    "index": 281474976710656,
```

```
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/3a2a8d78-ded3-11ed-9806-
005056ac3135/281474976710656"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume2",
      "uuid": "3a2a8d78-ded3-11ed-9806-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3a2a8d78-ded3-11ed-9806-
005056ac3135"
        }
      }
    },
    "index": 1153202979583557632,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/3a2a8d78-ded3-11ed-9806-
005056ac3135/1153202979583557632"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
```

```
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume2",
      "uuid": "3a2a8d78-ded3-11ed-9806-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3a2a8d78-ded3-11ed-9806-
005056ac3135"
        }
      }
    },
    "index": 3459045988797251584,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/3a2a8d78-ded3-11ed-9806-
005056ac3135/3459045988797251584"
      }
    }
  },
  {
    "svm": {
      "name": "vs0",
      "uuid": "842981cb-c985-11ed-a399-005056ac442b",
      "_links": {
        "self": {
          "href": "/api/svm/svms/842981cb-c985-11ed-a399-005056ac442b"
        }
      }
    },
    "volume": {
      "name": "srcVolume2",
      "uuid": "3a2a8d78-ded3-11ed-9806-005056ac3135",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3a2a8d78-ded3-11ed-9806-
005056ac3135"
        }
      }
    },
```

```
    "index": 4611967493404098560,
    "qtree": {
      "name": "qt1"
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/3a2a8d78-ded3-11ed-9806-
005056ac3135/4611967493404098560"
      }
    }
  }
],
"num_records": 9,
"_links": {
  "self": {
    "href":
"/api/storage/quota/reports?qtree.name=qt1&files.hard_limit=>5&#124;null"
  }
}
}
----

'''
```


```
[[IDa761c2905a26254c23c5894ad8ffadd0]]
= Retrieve the quota report records for all FlexVol and FlexGroup volumes

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/quota/reports`#

*Introduced In:* 9.6

Retrieves the quota report records for all FlexVol volumes and FlexGroup
volumes.

== Related ONTAP commands

* `quota report`


== Parameters

[cols=5*,options=header]
```

```
|===

|Name
|Type
|In
|Required
|Description

|space.soft_limit
|integer
|query
|False
a|Filter by space.soft_limit


|space.used.soft_limit_percent
|integer
|query
|False
a|Filter by space.used.soft_limit_percent


|space.used.total
|integer
|query
|False
a|Filter by space.used.total


|space.used.hard_limit_percent
|integer
|query
|False
a|Filter by space.used.hard_limit_percent


|space.hard_limit
|integer
|query
|False
a|Filter by space.hard_limit


|index
|integer
|query
|False
```

```
a|Filter by index


|files.soft_limit
|integer
|query
|False
a|Filter by files.soft_limit


|files.hard_limit
|integer
|query
|False
a|Filter by files.hard_limit


|files.used.total
|integer
|query
|False
a|Filter by files.used.total


|files.used.hard_limit_percent
|integer
|query
|False
a|Filter by files.used.hard_limit_percent


|files.used.soft_limit_percent
|integer
|query
|False
a|Filter by files.used.soft_limit_percent


|volume.name
|string
|query
|False
a|Filter by volume.name


|volume.uuid
|string
```

```
|query
|False
a|Filter by volume.uuid


|qtree.id
|integer
|query
|False
a|Filter by qtree.id


|qtree.name
|string
|query
|False
a|Filter by qtree.name


|users.id
|string
|query
|False
a|Filter by users.id


|users.name
|string
|query
|False
a|Filter by users.name


|type
|string
|query
|False
a|Filter by type


|group.id
|string
|query
|False
a|Filter by group.id
```

```
|group.name
|string
|query
|False
a|Filter by group.name


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|show_default_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, the default
records are not reported.

* Introduced in: 9.7
* Default value: 1


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
```

```
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#quota_report[quota_report]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "files": {
        "hard_limit": 0,
        "soft_limit": 0,
        "used": {
          "hard_limit_percent": 0,
          "soft_limit_percent": 0,
          "total": 0
        }
      },
      "group": {
        "id": "string",
        "name": "string"
      },
```

```
      "index": 0,
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 1,
        "name": "qt1"
      },
      "space": {
        "hard_limit": 0,
        "soft_limit": 0,
        "used": {
          "hard_limit_percent": 0,
          "soft_limit_percent": 0,
          "total": 0
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "users": [
        {
          "id": "string",
          "name": "string"
        }
      ],
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
```

```
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|

|===


[#used]
[.api-collapsible-fifth-title]
used

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit_percent
|integer
a|Total files used as a percentage of file hard limit


|soft_limit_percent
|integer
a|Total files used as a percentage of file soft limit


|total
|integer
a|Total files used


|===


[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
```

```
|integer
a|File hard limit


|soft_limit
|integer
a|File soft limit


|used
|link:#used[used]
a|

|===


[#group]
[.api-collapsible-fifth-title]
group

This parameter specifies the target group associated with the given quota
report record. This parameter is available for group quota records and is
not available for user or tree quota records. The target group is
identified by a UNIX group name and UNIX group identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target group ID


|name
|string
a|Quota target group name


|===


[#qtree]
[.api-collapsible-fifth-title]
```

```
qtree

This parameter specifies the target qtree associated with the user, group,
or tree record. For a user/group quota policy rule at volume level, this
parameter is not valid. For a default tree quota policy rule, this
parameter is specified as "" or "*". For a tree quota policy rule at qtree
level, this parameter specifies a qtree name and a qtree identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.


|===


[#used]
[.api-collapsible-fifth-title]
used

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit_percent
|integer
a|Total space used as a percentage of space hard limit
```

```
|soft_limit_percent
|integer
a|Total space used as a percentage of space soft limit


|total
|integer
a|Total space used


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|Space hard limit in bytes


|soft_limit
|integer
a|Space soft limit in bytes


|used
|link:#used[used]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID


|name
|string
a|Quota target user name


|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#quota_report]
[.api-collapsible-fifth-title]
quota_report

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|files
|link:#files[files]
a|

|group
|link:#group[group]
a|This parameter specifies the target group associated with the given
quota report record. This parameter is available for group quota records
and is not available for user or tree quota records. The target group is
identified by a UNIX group name and UNIX group identifier.


|index
|integer
a|Index that identifies a unique quota record. Valid in URL.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree associated with the user,
group, or tree record. For a user/group quota policy rule at volume level,
this parameter is not valid. For a default tree quota policy rule, this
parameter is specified as "" or "*". For a tree quota policy rule at qtree
level, this parameter specifies a qtree name and a qtree identifier.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Quota type associated with the quota record.


|users
|array[link:#users[users]]
a|This parameter specifies the target user or users associated with the
given quota report record. This parameter is available for user quota
records and is not available for group or tree quota records. The target
user or users are identified by a user name and user identifier. The user
```

```
name can be a UNIX user name or a Windows user name, and the identifier
can be a UNIX user identifier or a Windows security identifier.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDfd1cf7a52feed6bef71f423ef68b4f92]]
= Retrieve a specific quota report record

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/quota/reports/{volume.uuid}/{index}`#

*Introduced In:* 9.6

Retrieves a specific quota report record.

== Related ONTAP commands

* `quota report`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID


|index
|integer
|path
|True
a|Quota report index


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|files
|link:#files[files]
a|

|group
|link:#group[group]
```

```
a|This parameter specifies the target group associated with the given
quota report record. This parameter is available for group quota records
and is not available for user or tree quota records. The target group is
identified by a UNIX group name and UNIX group identifier.


|index
|integer
a|Index that identifies a unique quota record. Valid in URL.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree associated with the user,
group, or tree record. For a user/group quota policy rule at volume level,
this parameter is not valid. For a default tree quota policy rule, this
parameter is specified as "" or "*". For a tree quota policy rule at qtree
level, this parameter specifies a qtree name and a qtree identifier.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|Quota type associated with the quota record.


|users
|array[link:#users[users]]
a|This parameter specifies the target user or users associated with the
given quota report record. This parameter is available for user quota
records and is not available for group or tree quota records. The target
user or users are identified by a user name and user identifier. The user
name can be a UNIX user name or a Windows user name, and the identifier
can be a UNIX user identifier or a Windows security identifier.


|volume
|link:#volume[volume]
a|
```

```
|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "files": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "index": 0,
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 1,
    "name": "qt1"
  },
  "space": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
```

```
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```


Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 918235
| A volume with UUID was not found.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#used]
[.api-collapsible-fifth-title]
used

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit_percent
|integer
a|Total files used as a percentage of file hard limit


|soft_limit_percent
|integer
a|Total files used as a percentage of file soft limit
```

```
|total
|integer
a|Total files used


|===


[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|File hard limit


|soft_limit
|integer
a|File soft limit


|used
|link:#used[used]
a|

|===


[#group]
[.api-collapsible-fifth-title]
group

This parameter specifies the target group associated with the given quota
report record. This parameter is available for group quota records and is
not available for user or tree quota records. The target group is
identified by a UNIX group name and UNIX group identifier.


[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|id
|string
a|Quota target group ID


|name
|string
a|Quota target group name


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

This parameter specifies the target qtree associated with the user, group,
or tree record. For a user/group quota policy rule at volume level, this
parameter is not valid. For a default tree quota policy rule, this
parameter is specified as "" or "*". For a tree quota policy rule at qtree
level, this parameter specifies a qtree name and a qtree identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.
```

```
|===


[#used]
[.api-collapsible-fifth-title]
used

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit_percent
|integer
a|Total space used as a percentage of space hard limit


|soft_limit_percent
|integer
a|Total space used as a percentage of space soft limit


|total
|integer
a|Total space used


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|Space hard limit in bytes
```

```
|soft_limit
|integer
a|Space soft limit in bytes


|used
|link:#used[used]
a|

|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID


|name
|string
a|Quota target user name


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
```

```
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage quota policy rules

:leveloffset: +1


[[IDbd4fa9c92bac987c210ac18552eb7484]]
= Manage storage quota policy rules



== Overview

Quotas are defined in quota rules specific to FlexVol volumes or FlexGroup
volumes.  Each quota rule has a type. The type can be "user", "group", or
"tree".

&ndash; User rules must have the user property and qtree property.

&ndash; Group rules must have the group property and qtree property.

&ndash; Tree rules must have the qtree property and not have the user or
group property.
```

```
== Quota policy rule APIs

The following APIs can be used to perform create, retrieve, modify, and
delete operations related to quota policy rules.

&ndash; POST      /api/storage/quota/rules

&ndash; GET       /api/storage/quota/rules

&ndash; GET       /api/storage/quota/rules/\{rule-uuid}

&ndash; PATCH     /api/storage/quota/rules/\{rule-uuid}

&ndash; DELETE    /api/storage/quota/rules/\{rule-uuid}

== Examples

=== Retrieving all quota policy rules

This API is used to retrieve all quota policy rules.

The following example shows how to retrieve quota policy rules for FlexVol
volumes and FlexGroup volumes.

'''

----

# The API:
GET /api/storage/quota/rules

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules' -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
          }
        }
```

```
      },
      "volume": {
        "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-
4f142fc8cc92"
          }
        }
      },
      "uuid": "66319cbe-b837-11e8-9c5a-005056a7e88c",
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/66319cbe-b837-11e8-9c5a-
005056a7e88c"
        }
      }
    },
    {
      "svm": {
        "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
          }
        }
      },
      "volume": {
        "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-
4f142fc8cc92"
          }
        }
      },
      "uuid": "dbd5b443-b7a4-11e8-bc58-005056a7e88c",
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/dbd5b443-b7a4-11e8-bc58-
005056a7e88c"
        }
      }
```

```
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules"
    }
  }
}
----

'''

=== Retrieving a specific quota policy rule

This API is used to retrieve a quota policy rule for a specific qtree.

The following example shows how to retrieve a quota policy user rule for a
specific qtree.

'''

----

# The API:
GET /api/storage/quota/rules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/264a9e0b-2e03-11e9-
a610-005056a7b72d' -H 'accept: application/hal+json'

# Response for a user rule at a qtree level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
```

```
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "264a9e0b-2e03-11e9-a610-005056a7b72d",
  "type": "user",
  "users": [ {"name" : "fred"} ],
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    }
  },
  "user_mapping": on,
  "space": {
    "hard_limit": 1222800,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 80
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/264a9e0b-2e03-11e9-a610-
005056a7b72d"
    }
  }
}
----

'''

=== Retrieving a quota policy multi-user rule at the volume level

'''

----

# The call:
```

```
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/0ab84fba-19aa-11e9-
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a multi-user rule at volume level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "0ab84fba-19aa-11e9-a04d-005056a72f42",
  "type": "user",
  "users": [
    {
      "name": "sam",
    },
    {
      "name": "smith",
    },
    {
      "id": "300010",
    },
  ],
  "space": {
    "hard_limit": 1222800,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 80
  },
  "_links": {
```

```
    "self": {
      "href": "/api/storage/quota/rules/0ab84fba-19aa-11e9-a04d-
005056a72f42"
    }
  }
}
----

'''

=== Retrieving a quota policy default tree rule

'''

----

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/4a276b8c-1753-11e9-
8101-005056a760e0' -H 'accept: application/hal+json'

# Response for a default tree rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "4a276b8c-1753-11e9-8101-005056a760e0",
  "type": "tree",
  "qtree": {
    "name": ""
  },
```

```
  "space": {
    "hard_limit": 1034000,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/4a276b8c-1753-11e9-8101-
005056a760e0"
    }
  }
}
----

'''

=== Retrieving a quota policy tree rule for a specific qtree

'''

----

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/49b1134f-19ab-11e9-
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a tree rule for a specific qtree:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
```

```
ab489830d268"
      }
    }
  },
  "uuid": "49b1134f-19ab-11e9-a04d-005056a72f42",
  "type": "tree",
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    }
  },
  "space": {
    "hard_limit": 1048576,
    "soft_limit": 838861
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 40
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/49b1134f-19ab-11e9-a04d-
005056a72f42"
    }
  }
}
----

'''

=== Retrieving a quota policy group rule for a specific qtree

'''

----

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/b9236852-19ab-11e9-
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a group rule:
```

```
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "b9236852-19ab-11e9-a04d-005056a72f42",
  "type": "group",
  "group": {"name" : "group1"},
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    }
  },
  "space": {
    "hard_limit": 2097152,
    "soft_limit": 1572864
  },
  "files": {
    "hard_limit": 250,
    "soft_limit": 200
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/b9236852-19ab-11e9-a04d-
005056a72f42"
    }
```

```
  }
}
----

'''

=== Creating a quota policy rule

This API is used to create a new quota policy rule. When an explicit rule
or a qtree-scoped rule of a type is created on a volume, a default rule of
the same type is automatically added if it does not already exist on the
volume.

The following example shows how to create a quota policy user rule using
POST.

'''

----

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "vol1"
},
"type": "user",
"users": [ {"name" : "example_user"} ],
"qtree": {
  "name":"qt1"
},
"user_mapping": "on",
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
```

```
  "hard_limit": 20,
  "soft_limit": 10
}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3220eea6-5049-11e9-bfb7-005056a7f717",
      "type": "user",
      "users": [
        {
          "name" : "example_user"
        }
      ],
      "qtree": {
        "name": "qt1"
      },
      "user_mapping": "on",
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/3220eea6-5049-11e9-bfb7-
005056a7f717"
        }
      }
    }
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
```

```
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}
----

'''

=== Creating a quota policy group rule using POST.

'''

----

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "vol1"
},
"type": "group",
"group": {
  "name" : "test_group1"
},
"qtree": {
  "name":"qt1"
},
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
```

```
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3b130f7d-504a-11e9-bfb7-005056a7f717",
      "type": "group",
      "group": {
          "name" : "test_group1"
      },
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/3b130f7d-504a-11e9-bfb7-
005056a7f717"
        }
      }
    }
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}
```

```
----

'''

=== Creating a quota policy tree rule using POST

'''

----

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "vol1"
},
"type": "tree",
"qtree": {
  "name":"qt1"
},
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
```

```
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "e5eb03be-504a-11e9-bfb7-005056a7f717",
      "type": "tree",
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/e5eb03be-504a-11e9-bfb7-
005056a7f717"
        }
      }
    }
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}
----

'''

=== Updating the quota policy rule

This API is used to update a quota policy rule.

The following example shows how to update a quota policy rule.

'''
```

```
----

# The API:
PATCH /storage/quota/rules/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-
11e8-a806-005056a7e73a" -H 'accept: application/hal+json' -d
"@test_quota_patch.txt"
test_quota_patch.txt(body):
{
"space": {
  "hard_limit": 16554,
  "soft_limit": 8192
},
"files": {
  "hard_limit": 40,
  "soft_limit": 20
}
}
----

'''

=== Deleting the quota policy rule

This API is used to delete a quota policy rule.

The following example shows how to delete a quota policy rule.

'''

----

# The API:
DELETE /storage/quota/rules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-
11e8-a806-005056a7e73a" -H 'accept: application/hal+json'
----

'''

=== Retrieving a quota policy rule with the property files.hard_limit
greater than 5 or null (unlimited) for qtree qt1
```

```
'''

----

# The API:
GET /api/storage/quota/rules

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/rules?rqtree.name=qt1&files.hard_limit=>5&#124;null"
-H 'accept: application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "2272378d-eab2-11ed-913b-005056ac10bf",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2272378d-eab2-11ed-913b-005056ac10bf"
        }
      }
    },
    "volume": {
      "uuid": "6dc24682-edde-11ed-8b6b-005056acde94",
      "name": "testVol",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/6dc24682-edde-11ed-8b6b-
005056acde94"
        }
      }
    },
    "uuid": "9abfdd9f-ede7-11ed-8b6b-005056acde94",
    "files": {
      "hard_limit": 15
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/9abfdd9f-ede7-11ed-8b6b-
005056acde94"
      }
    }
  },
```

```
  {
    "svm": {
      "uuid": "2272378d-eab2-11ed-913b-005056ac10bf",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2272378d-eab2-11ed-913b-005056ac10bf"
        }
      }
    },
    "volume": {
      "uuid": "83c865bd-edde-11ed-8b6b-005056acde94",
      "name": "testVol2",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/83c865bd-edde-11ed-8b6b-
005056acde94"
        }
      }
    },
    "uuid": "a876601e-ede7-11ed-8b6b-005056acde94",
    "files": {
      "hard_limit": 6
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/a876601e-ede7-11ed-8b6b-
005056acde94"
      }
    }
  },
  {
    "svm": {
      "uuid": "2272378d-eab2-11ed-913b-005056ac10bf",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2272378d-eab2-11ed-913b-005056ac10bf"
        }
      }
    },
    "volume": {
      "uuid": "a44c8f5d-edde-11ed-8b6b-005056acde94",
      "name": "testVol4",
      "_links": {
        "self": {
```

```
          "href": "/api/storage/volumes/a44c8f5d-edde-11ed-8b6b-
005056acde94"
        }
      }
    },
    "uuid": "d66ff5ed-ede7-11ed-8b6b-005056acde94",
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/d66ff5ed-ede7-11ed-8b6b-
005056acde94"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/quota/rules?files.hard_limit=>5&#124;null"
  }
}
}
----

'''
```

```
[[IDaa514f6e41e92977cb440c326433bc42]]
= Retrieve quota policy rules for all FlexVol and FlexGroup volumes

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/quota/rules`#

*Introduced In:* 9.6

Retrieves quota policy rules configured for all FlexVol volumes and
FlexGroup volumes.

== Related ONTAP commands

* `quota policy rule show`


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|volume.name
|string
|query
|False
a|Filter by volume.name


|volume.uuid
|string
|query
|False
a|Filter by volume.uuid


|qtree.name
|string
|query
|False
a|Filter by qtree.name


|qtree.id
|integer
|query
|False
a|Filter by qtree.id


|files.hard_limit
|integer
|query
|False
a|Filter by files.hard_limit


|files.soft_limit
|integer
|query
```

```
|False
a|Filter by files.soft_limit


|space.hard_limit
|integer
|query
|False
a|Filter by space.hard_limit


|space.soft_limit
|integer
|query
|False
a|Filter by space.soft_limit


|uuid
|string
|query
|False
a|Filter by uuid


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|user_mapping
|boolean
|query
|False
a|Filter by user_mapping


|group.name
```

```
|string
|query
|False
a|Filter by group.name


|group.id
|string
|query
|False
a|Filter by group.id


|type
|string
|query
|False
a|Filter by type


|users.id
|string
|query
|False
a|Filter by users.id


|users.name
|string
|query
|False
a|Filter by users.name


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|num_records
|integer
a|Number of records


|records
|array[link:#quota_rule[quota_rule]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "group": {
        "id": "string",
        "name": "string"
      },
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 1,
```

```
        "name": "qt1"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "users": [
        {
          "id": "string",
          "name": "string"
        }
      ],
      "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
```

```
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|hard_limit
|integer
a|This parameter specifies the hard limit for files. This is valid in POST
or PATCH.


|soft_limit
|integer
a|This parameter specifies the soft limit for files. This is valid in POST
or PATCH.


|===


[#group]
[.api-collapsible-fifth-title]
group

This parameter specifies the target group to which the group quota policy
rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target group ID


|name
|string
a|Quota target group name
```

```
|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the space hard limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|soft_limit
|integer
a|This parameter specifies the space soft limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID


|name
|string
a|Quota target user name


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#quota_rule]
[.api-collapsible-fifth-title]
quota_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|files
|link:#files[files]
a|

|group
|link:#group[group]
a|This parameter specifies the target group to which the group quota
policy rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
```

```
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|This parameter specifies the quota policy rule type. This is required in
POST only and can take either one of the "user", "group" or "tree" values.


|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
valid in POST or PATCH for user quota policy rules only.


|users
|array[link:#users[users]]
a|This parameter specifies the target user to which the user quota policy
rule applies. This parameter takes single or multiple user names or
identifiers. This parameter is valid only for the POST operation of a user
quota policy rule. If this parameter is used as an input to create a group
or a tree quota policy rule, the POST operation will fail with an
appropriate error. For POST, this input parameter takes either a user name
or a user identifier, not both. For default quota rules, the user name
```

```
must be chosen and specified as "". For explicit user quota rules, this
parameter can indicate either a user name or user identifier. The user
name can be a UNIX user name or a Windows user name. If a name contains a
space, enclose the entire value in quotes. A UNIX user name cannot include
a backslash () or an @ sign; user names with these characters are treated
as Windows names. The user identifier can be a UNIX user identifier or a
Windows security identifier. For multi-user quota, this parameter can
contain multiple user targets separated by a comma.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|volume
|link:#volume[volume]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDfd4bf61aa9d78823e97a91d054b4f805]]
= Create a quota policy rule for a FlexVol or a FlexGroup volume

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/quota/rules`#

*Introduced In:* 9.6

Creates a quota policy rule for a FlexVol or a FlexGroup volume.
```

```
Important notes:

* Unlike CLI/ONTAPI, the `quota policy` input is not needed for POST.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the qtree.
* `volume.uuid` or `volume.name` - Existing volume in which to create the
qtree.
* `type` - Quota type for the rule. This type can be `user`, `group`, or
`tree`.
* `users.name` or `user.id` -  If the quota type is user, this property
takes the user name or user ID. For default user quota rules, the user
name must be specified as "".
* `group.name` or `group.id` - If the quota type is group, this property
takes the group name or group ID. For default group quota rules, the group
name must be specified as "".
* `qtree.name` - Qtree for which to create the rule. For default tree
rules, the qtree name must be specified as "".

== Recommended optional properties

* `space.hard_limit` - Specifies the space hard limit, in bytes. If less
than 1024 bytes, the value is rounded up to 1024 bytes.
* `space.soft_limit` - Specifies the space soft limit, in bytes. If less
than 1024 bytes, the value is rounded up to 1024 bytes.
* `files.hard_limit` - Specifies the hard limit for files.
* `files.soft_limit` - Specifies the soft limit for files.
* `user_mapping` - Specifies the user_mapping. This property is valid only
for quota policy rules of type `user`.

== Related ONTAP commands

* `quota policy rule create`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|files
|link:#files[files]
a|

|group
|link:#group[group]
a|This parameter specifies the target group to which the group quota
policy rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
```

```
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|This parameter specifies the quota policy rule type. This is required in
POST only and can take either one of the "user", "group" or "tree" values.


|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
valid in POST or PATCH for user quota policy rules only.


|users
|array[link:#users[users]]
a|This parameter specifies the target user to which the user quota policy
rule applies. This parameter takes single or multiple user names or
identifiers. This parameter is valid only for the POST operation of a user
quota policy rule. If this parameter is used as an input to create a group
```

```
or a tree quota policy rule, the POST operation will fail with an
appropriate error. For POST, this input parameter takes either a user name
or a user identifier, not both. For default quota rules, the user name
must be chosen and specified as "". For explicit user quota rules, this
parameter can indicate either a user name or user identifier. The user
name can be a UNIX user name or a Windows user name. If a name contains a
space, enclose the entire value in quotes. A UNIX user name cannot include
a backslash () or an @ sign; user names with these characters are treated
as Windows names. The user identifier can be a UNIX user identifier or a
Windows security identifier. For multi-user quota, this parameter can
contain multiple user targets separated by a comma.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|volume
|link:#volume[volume]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "group": {
    "id": "string",
    "name": "string"
  },
  "qtree": {
    "id": 1,
    "name": "qt1"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
```

```
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 917927
| The specified volume was not found.

| 918232
| Either `volume.name` or `volume.uuid` must be provided.

| 918236
| The specified `volume.uuid` and `volume.name` refer to different
volumes.

| 2621462
| The specified SVM does not exist.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5308469
| User mapping can be specified only for a user quota rule.

| 5308496
| The specified user name was not found.

| 5308501
| Mapping from Windows user to UNIX user for user rule was unsuccessful.

| 5308502
| Mapping from UNIX user to Windows user for user rule was unsuccessful.

| 5308552
| Failed to get default quota policy name for SVM.

| 5308561
| Failed to obtain volume quota state or invalid quota state obtained for
```

```
volume.

| 5308562
| users is a required input for creating a user rule and group is not
allowed.

| 5308563
| group is a required input for creating a group rule and users is not
allowed.

| 5308564
| qtree.name is a required input for creating a tree rule and users and
group are not allowed.

| 5308565
| Only one of name or id is allowed for each entry in the users array.

| 5308566
| Only one of name or id is allowed for group.

| 5308568
| Quota policy rule create operation succeeded, but quota resize failed
due to internal error. To activate the rule, disable and enable quotas for
this volume.

| 5308571
| Quota policy rule create operation succeeded, but quota resize is
skipped. To activate the rule, disable and enable quotas for this volume.

| 5308573
| Input value is greater than limit for field.

| 5308574
| Input value is out of range for field.

| 5308575
| Input value is incorrectly larger than listed field.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the hard limit for files. This is valid in POST
or PATCH.


|soft_limit
|integer
a|This parameter specifies the soft limit for files. This is valid in POST
or PATCH.


|===
```

```
[#group]
[.api-collapsible-fifth-title]
group

This parameter specifies the target group to which the group quota policy
rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target group ID


|name
|string
a|Quota target group name


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
```

```
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the space hard limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|soft_limit
|integer
a|This parameter specifies the space soft limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
```

```
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID
```

```
|name
|string
a|Quota target user name


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#quota_rule]
[.api-collapsible-fifth-title]
quota_rule

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|files
|link:#files[files]
a|

|group
|link:#group[group]
a|This parameter specifies the target group to which the group quota
policy rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|This parameter specifies the quota policy rule type. This is required in
POST only and can take either one of the "user", "group" or "tree" values.
```

```
|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
valid in POST or PATCH for user quota policy rules only.


|users
|array[link:#users[users]]
a|This parameter specifies the target user to which the user quota policy
rule applies. This parameter takes single or multiple user names or
identifiers. This parameter is valid only for the POST operation of a user
quota policy rule. If this parameter is used as an input to create a group
or a tree quota policy rule, the POST operation will fail with an
appropriate error. For POST, this input parameter takes either a user name
or a user identifier, not both. For default quota rules, the user name
must be chosen and specified as "". For explicit user quota rules, this
parameter can indicate either a user name or user identifier. The user
name can be a UNIX user name or a Windows user name. If a name contains a
space, enclose the entire value in quotes. A UNIX user name cannot include
a backslash () or an @ sign; user names with these characters are treated
as Windows names. The user identifier can be a UNIX user identifier or a
Windows security identifier. For multi-user quota, this parameter can
contain multiple user targets separated by a comma.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|volume
|link:#volume[volume]
a|

|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDd2b1acc794d36c1b42d0bbbf63894e1d]]
= Delete a quota policy rule

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/quota/rules/{uuid}`#

*Introduced In:* 9.7

Deletes a quota policy rule.

== Related ONTAP commands

* `quota policy rule delete`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
```

```
|In
|Required
|Description

|uuid
|string
|path
|True
a|Rule UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5308545
| The specified quota rule UUID is invalid.
```

```
| 5308561
| Failed to obtain volume quota state or invalid quota state obtained for
volume.

| 5308569
| Quota policy rule delete operation succeeded, but quota resize failed
due to internal error.

| 5308572
| Quota policy rule delete operation succeeded, however the rule is still
being enforced. To stop enforcing the rule, disable quotas and enable them
again for this volume.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
```

```
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf535e6765c8ad6a99698f57e2430de83]]
= Retrieve quota policy rule properties

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/quota/rules/{uuid}`#

*Introduced In:* 9.7

Retrieves properties for a specific quota policy rule.
```

```
== Related ONTAP commands

* `quota policy rule show`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Rule UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|files
```

```
|link:#files[files]
a|

|group
|link:#group[group]
a|This parameter specifies the target group to which the group quota
policy rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


|qtree
|link:#qtree[qtree]
a|This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


|space
|link:#space[space]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|type
|string
a|This parameter specifies the quota policy rule type. This is required in
POST only and can take either one of the "user", "group" or "tree" values.


|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
```

```
valid in POST or PATCH for user quota policy rules only.


|users
|array[link:#users[users]]
a|This parameter specifies the target user to which the user quota policy
rule applies. This parameter takes single or multiple user names or
identifiers. This parameter is valid only for the POST operation of a user
quota policy rule. If this parameter is used as an input to create a group
or a tree quota policy rule, the POST operation will fail with an
appropriate error. For POST, this input parameter takes either a user name
or a user identifier, not both. For default quota rules, the user name
must be chosen and specified as "". For explicit user quota rules, this
parameter can indicate either a user name or user identifier. The user
name can be a UNIX user name or a Windows user name. If a name contains a
space, enclose the entire value in quotes. A UNIX user name cannot include
a backslash () or an @ sign; user names with these characters are treated
as Windows names. The user identifier can be a UNIX user identifier or a
Windows security identifier. For multi-user quota, this parameter can
contain multiple user targets separated by a comma.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|volume
|link:#volume[volume]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "group": {
```

```
    "id": "string",
    "name": "string"
  },
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 1,
    "name": "qt1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```


Status: Default

```
```

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 5308544
| The specified quota rule UUID is invalid.

| 5308545
| Unable to retrieve rule for the specified quota rule UUID.

| 5308576
| Parameter `show_default_records` only allowed for GET collection.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===
```

```
[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the hard limit for files. This is valid in POST
or PATCH.


|soft_limit
|integer
a|This parameter specifies the soft limit for files. This is valid in POST
or PATCH.


|===


[#group]
[.api-collapsible-fifth-title]
group

This parameter specifies the target group to which the group quota policy
rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
```

```
|string
a|Quota target group ID


|name
|string
a|Quota target group name


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.
```

```
|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the space hard limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|soft_limit
|integer
a|This parameter specifies the space soft limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID


|name
|string
a|Quota target user name


|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDf1797163da5de48beec6cad087a7124c]]
= Update quota policy rule properties

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/quota/rules/{uuid}`#
```

```
*Introduced In:* 9.7

Updates properties of a specific quota policy rule.

Important notes:

* The quota resize functionality is supported with the PATCH operation.
* Quota resize allows you to modify the quota limits, directly in the
filesystem.
* The quota must be enabled on a FlexVol or a FlexGroup volume for `quota
resize` to take effect.
* If the quota is disabled on the volume, the quota policy rule PATCH API
modifies the rule, but this does not affect the limits in the filesystem.

== Related ONTAP commands

* `quota policy rule modify`
* `quota resize`


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Rule UUID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
```

```
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|files
|link:#files[files]
a|

|space
|link:#space[space]
a|

|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
valid in POST or PATCH for user quota policy rules only.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
```

```
| Error Code | Description

| 5308478
| User mapping can only be turned on for UNIX usernames or Windows account
names.

| 5308501
| Mapping from Windows user to UNIX user for user rule was unsuccessful.

| 5308502
| Mapping from UNIX user to Windows user for user rule was unsuccessful.

| 5308545
| The specified quota rule UUID is invalid.

| 5308561
| Failed to obtain volume quota state or invalid quota state obtained for
volume.

| 5308567
| Quota policy rule modify operation succeeded, but quota resize failed
due to internal error.

| 5308573
| Input value is greater than limit for field.

| 5308574
| Input value is out of range for field.

| 5308575
| Input value is incorrectly larger than listed field.
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#files]
[.api-collapsible-fifth-title]
files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the hard limit for files. This is valid in POST
or PATCH.


|soft_limit
|integer
a|This parameter specifies the soft limit for files. This is valid in POST
or PATCH.


|===


[#group]
[.api-collapsible-fifth-title]
group
```

```
This parameter specifies the target group to which the group quota policy
rule applies. This parameter takes a group name or identifier. This
parameter is only valid for the POST operation of a group quota policy
rule. The POST operation will fail with an appropriate error if this
parameter is used as an input to create a user or a tree quota policy
rule. This input parameter for POST takes either a group name or a group
identifier, but not both. For default quota rules, the group name must be
chosen and should be specified as "". For explicit group quota rules, this
parameter can contain a UNIX group name or a UNIX group identifier.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target group ID


|name
|string
a|Quota target group name


|===


[#qtree]
[.api-collapsible-fifth-title]
qtree

This parameter specifies the target qtree to which the user/group/tree
quota policy rule applies. For a user/group quota policy rule at qtree
level, this parameter takes a qtree name and is valid in GET or POST. For
a user/group quota policy rule at volume level, this parameter is not
valid in GET or POST. For a tree quota policy rule, this parameter is
mandatory and is valid in both POST and GET. For a default tree quota
policy rule, this parameter needs to be specified as "". For a tree quota
policy rule at qtree level, this parameter takes a qtree name and is valid
in GET or POST.


[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|id
|integer
a|The unique identifier for a qtree.


|name
|string
a|The name of the qtree.


|===


[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|hard_limit
|integer
a|This parameter specifies the space hard limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|soft_limit
|integer
a|This parameter specifies the space soft limit, in bytes. If less than
1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
For a POST operation where the parameter is either empty or set to -1, no
limit is applied. For a PATCH operation where a limit is configured, use a
value of -1 to clear the limit.


|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#users]
[.api-collapsible-fifth-title]
users

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|Quota target user ID


|name
|string
a|Quota target user name
```

```
|===


[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the volume. This field cannot be specified in a PATCH
method.


|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
* x-nullable: true


|===


[#quota_rule]
[.api-collapsible-fifth-title]
quota_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|files
|link:#files[files]
a|
```

```
|space
|link:#space[space]
a|

|user_mapping
|boolean
a|This parameter enables user mapping for user quota policy rules. This is
valid in POST or PATCH for user quota policy rules only.


|uuid
|string
a|Unique identifier for the quota policy rule. This field is generated
when the quota policy rule is created.


|===


[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage shelves

:leveloffset: +1


[[IDf2fc6055ebae6c8bb4c272b9760ad490]]
= Manage shelves



== Retrieving storage shelf information

The storage shelf GET API retrieves all of the shelves in the cluster.

'''

== Examples

=== 1) Retrieve a list of shelves from the cluster

==== The following example shows the response with a list of shelves in
the cluster:

'''

----

# The API:
/api/storage/shelves

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/shelves" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
```

```
    "uid": "3109174803597886800",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/3109174803597886800"
      }
    }
  },
  {
    "uid": "9237728366621690448",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/9237728366621690448"
      }
    }
  },
  {
    "uid": "9946762738829886800",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/9946762738829886800"
      }
    }
  },
  {
    "uid": "10318311901725526608",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/10318311901725526608"
      }
    }
  },
  {
    "uid": "13477584846688355664",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/13477584846688355664"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/shelves/"
  }
}
```

```
}
----

'''

=== 2) Retrieve a specific shelf from the cluster

==== The following example shows the response of the requested shelf. If
there is no shelf with the requested uid, an error is returned.

'''

----

# The API:
/api/storage/shelves/{uid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/shelves/3109174803597886800" -H
"accept: application/hal+json"

# The response:
{
"uid": "3109174803597886800",
"name": "6.10",
"id": "10",
"serial_number": "SHU0954292N0HAH",
"model": "DS4246",
"module_type": "iom6",
"internal": false,
"local": true,
"manufacturer": {
  "name": "NETAPP"
},
"state": "ok",
"connection_type": "sas",
"disk_count": 24,
"raw-capacity": 5157308772,
"location_led": "off",
"paths": [
  {
    "name": "0e",
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
```

```
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0e"
      }
    }
  },
  {
    "name": "0g",
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0g"
      }
    }
  }
],
"bays": [
  {
    "id": 0,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 1,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
```

```
  {
    "id": 2,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 3,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 4,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 5,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 6,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 7,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 8,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 9,
    "has_disk": true,
    "type": "single_disk",
```

```
      "state": "ok"
    },
    {
      "id": 10,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 11,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 12,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 13,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 14,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 15,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 16,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 17,
```

```
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 18,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 19,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 20,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 21,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 22,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 23,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    }
  ],
  "frus": [
    {
      "type": "module",
      "id": 0,
      "state": "ok",
```

```
    "part_number": "111-00690+B2",
    "serial_number": "8001900099",
    "firmware_version": "0191",
    "installed": true
  },
  {
    "type": "module",
    "id": 1,
    "state": "ok",
    "part_number": "111-00190+B0",
    "serial_number": "7903785183",
    "firmware_version": "0191",
    "installed": true
  },
  {
    "type": "psu",
    "id": 1,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW82562007513E",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  },
  {
    "type": "psu",
    "id": 2,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW825620075138",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  },
  {
    "type": "psu",
    "id": 3,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW8256200750BA",
    "firmware_version": "0311",
    "installed": true,
```

```
      "psu": {
        "model": "9C"
      }
    },
    {
      "type": "psu",
      "id": 4,
      "state": "ok",
      "part_number": "0082562-12",
      "serial_number": "PMW8256200750A2",
      "firmware_version": "0311",
      "installed": true,
      "psu": {
        "model": "9C"
      }
    }
  ],
  "ports": [
    {
      "id": 0,
      "module_id": "a",
      "designator": "square",
      "state": "connected",
      "internal": false,
      "wwn": "500A098000C9EDBF",
      "speed": 100,
      "cable": {
        "identifier": "5001086000702488-500a098000c9edbf",
        "part_number": "112-00430+A0",
        "length": "2m",
        "serial_number": "APF16510229807"
      },
      "remote": {
        "wwn": "5001086000702488",
        "phy": "08"
      }
    },
    {
      "id": 1,
      "module_id": "a",
      "designator": "circle",
      "state": "connected",
      "internal": false,
      "wwn": "500A098000C9EDBF",
      "speed": 100,
      "cable": {
```

```
      "identifier": "500a098000d5c4bf-500a098000c9edbf",
      "part_number": "112-00176+A0",
      "length": "0.5-1.0m",
      "serial_number": "APF133917610YT"
    },
    "remote": {
      "wwn": "500A098000D5C4BF",
      "phy": "00"
    }
  },
  {
    "id": 2,
    "module_id": "b",
    "designator": "square",
    "state": "connected",
    "internal": false,
    "wwn": "500A098004F208BF",
    "speed": 100,
    "cable": {
      "identifier": "5001086000702648-500a098004f208bf",
      "part_number": "112-00430+A0",
      "length": "2m",
      "serial_number": "APF16510229540"
    },
    "remote": {
      "wwn": "5001086000702648",
      "phy": "08"
    }
  },
  {
    "id": 3,
    "module_id": "b",
    "designator": "circle",
    "state": "connected",
    "internal": false,
    "wwn": "500A098004F208BF",
    "speed": 100,
    "cable": {
      "identifier": "500a0980062ba33f-500a098004f208bf",
      "part_number": "112-00176+20",
      "length": "0.5-1.0m",
      "serial_number": "832210017"
    },
    "remote": {
      "wwn": "500A0980062BA33F",
      "phy": "00"
```

```
      }
    }
  ],
  "fans": [
    {
      "id": 1,
      "location": "rear of the shelf on the upper left power supply",
      "rpm": 3150,
      "state": "ok",
      "installed": true
    },
    {
      "id": 2,
      "location": "rear of the shelf on the upper left power supply",
      "rpm": 3000,
      "state": "ok",
      "installed": true
    },
    {
      "id": 3,
      "location": "rear of the shelf on the upper right power supply",
      "rpm": 3220,
      "state": "ok",
      "installed": true
    },
    {
      "id": 4,
      "location": "rear of the shelf on the upper right power supply",
      "rpm": 3000,
      "state": "ok",
      "installed": true
    },
    {
      "id": 5,
      "location": "rear of the shelf on the lower left power supply",
      "rpm": 3000,
      "state": "ok",
      "installed": true
    },
    {
      "id": 6,
      "location": "rear of the shelf on the lower left power supply",
      "rpm": 3150,
      "state": "ok",
      "installed": true
    },
```

```
    {
      "id": 7,
      "location": "rear of the shelf on the lower right power supply",
      "rpm": 3150,
      "state": "ok",
      "installed": true
    },
    {
      "id": 8,
      "location": "rear of the shelf on the lower right power supply",
      "rpm": 3000,
      "state": "ok",
      "installed": true
    }
  ],
  "temperature_sensors": [
    {
      "id": 1,
      "location": "front of the shelf on the left, on the OPS panel",
      "temperature": 20,
      "ambient": true,
      "state": "ok",
      "installed": true,
      "threshold": {
        "high": {
          "critical": 42,
          "warning": 40
        },
        "low": {
          "critical": 0,
          "warning": 5
        }
      }
    },
    {
      "id": 2,
      "location": "inside of the shelf on the midplane",
      "temperature": 29,
      "ambient": false,
      "state": "ok",
      "installed": true,
      "threshold": {
        "high": {
          "critical": 55,
          "warning": 50
        },
```

```
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 3,
    "location": "rear of the shelf on the upper left power supply",
    "temperature": 33,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 55,
        "warning": 50
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 4,
    "location": "rear of the shelf on the upper left power supply",
    "temperature": 41,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 70,
        "warning": 65
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 5,
    "location": "rear of the shelf on the upper right power supply",
    "temperature": 32,
```

```
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 55,
        "warning": 50
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 6,
    "location": "rear of the shelf on the upper right power supply",
    "temperature": 41,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 70,
        "warning": 65
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 7,
    "location": "rear of the shelf on the lower left power supply",
    "temperature": 34,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 55,
        "warning": 50
      },
      "low": {
        "critical": 5,
```

```
        "warning": 10
      }
    }
  },
  {
    "id": 8,
    "location": "rear of the shelf on the lower left power supply",
    "temperature": 45,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 70,
        "warning": 65
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 9,
    "location": "rear of the shelf on the lower right power supply",
    "temperature": 30,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 55,
        "warning": 50
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 10,
    "location": "rear of the shelf on the lower right power supply",
    "temperature": 40,
    "ambient": false,
    "state": "ok",
```

```
    "installed": true,
    "threshold": {
      "high": {
        "critical": 70,
        "warning": 65
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 11,
    "location": "rear of the shelf at the top left, on shelf module A",
    "temperature": 30,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 60,
        "warning": 55
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
    }
  },
  {
    "id": 12,
    "location": "rear of the shelf at the top right, on shelf module B",
    "temperature": 33,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 60,
        "warning": 55
      },
      "low": {
        "critical": 5,
        "warning": 10
      }
```

```
      }
    }
  ],
  "voltage_sensors": [
    {
      "id": 1,
      "location": "rear of the shelf on the upper left power supply",
      "voltage": 5.11,
      "state": "ok",
      "installed": true
    },
    {
      "id": 2,
      "location": "rear of the shelf on the upper left power supply",
      "voltage": 12.38,
      "state": "ok",
      "installed": true
    },
    {
      "id": 3,
      "location": "rear of the shelf on the upper right power supply",
      "voltage": 5.11,
      "state": "ok",
      "installed": true
    },
    {
      "id": 4,
      "location": "rear of the shelf on the upper right power supply",
      "voltage": 12.26,
      "state": "ok",
      "installed": true
    },
    {
      "id": 5,
      "location": "rear of the shelf on the lower left power supply",
      "voltage": 5.7,
      "state": "ok",
      "installed": true
    },
    {
      "id": 6,
      "location": "rear of the shelf on the lower left power supply",
      "voltage": 12.26,
      "state": "ok",
      "installed": true
    },
```

```
    {
      "id": 7,
      "location": "rear of the shelf on the lower right power supply",
      "voltage": 5.15,
      "state": "ok",
      "installed": true
    },
    {
      "id": 8,
      "location": "rear of the shelf on the lower right power supply",
      "voltage": 12.3,
      "state": "ok",
      "installed": true
    }
  ],
  "current_sensors": [
    {
      "id": 1,
      "location": "rear of the shelf on the upper left power supply",
      "current": 6990,
      "state": "ok",
      "installed": true
    },
    {
      "id": 2,
      "location": "rear of the shelf on the upper left power supply",
      "current": 5150,
      "state": "ok",
      "installed": true
    },
    {
      "id": 3,
      "location": "rear of the shelf on the upper right power supply",
      "current": 4600,
      "state": "ok",
      "installed": true
    },
    {
      "id": 4,
      "location": "rear of the shelf on the upper right power supply",
      "current": 4800,
      "state": "ok",
      "installed": true
    },
    {
      "id": 5,
```

```
    "location": "rear of the shelf on the lower left power supply",
    "current": 4140,
    "state": "ok",
    "installed": true
  },
  {
    "id": 6,
    "location": "rear of the shelf on the lower left power supply",
    "current": 7770,
    "state": "ok",
    "installed": true
  },
  {
    "id": 7,
    "location": "rear of the shelf on the lower right power supply",
    "current": 4140,
    "state": "ok",
    "installed": true
  },
  {
    "id": 8,
    "location": "rear of the shelf on the lower right power supply",
    "current": 4720,
    "state": "ok",
    "installed": true
  }
],
"acps": [
  {
    "enabled": true,
    "channel": "in_band",
    "connection_state": "active",
    "node": {
      "uuid": "cf62d23c-6100-11eb-9852-00a098fd725d",
      "name": "cat33-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/cf62d23c-6100-11eb-9852-
00a098fd725d"
        }
      }
    }
  },
  {
    "enabled": true,
    "channel": "in_band",
```

```
    "connection_state": "active",
    "node": {
      "uuid": "d0892dd7-6100-11eb-9cdb-d039ea010238",
      "name": "cat33-02",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/d0892dd7-6100-11eb-9cdb-
d039ea010238"
        }
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/storage/shelves/3109174803597886800"
  }
}
}
----

'''

== Modifying storage shelf

The storage shelf PATCH API modifies the shelf location LED.

'''

== Example

'''

----

# The API:
/api/storage/shelves

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/shelves/3109174803597886800"
-H "accept: application/hal+json" -H "Content-Type: application/hal+json"
-d '{"location_led": "on"}'

# The response:
{
}
----
```

```
'''




[[ID003a559d7d5b4ae2447c539ab9fad526]]
= Retrieve shelves

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/shelves`#

*Introduced In:* 9.6

Retrieves a collection of shelves.

== Related ONTAP commands

* `storage shelf show`
* `storage shelf port show`
* `storage shelf drawer show`
* `storage shelf drawer show-slot`
* `storage shelf acp show`

== Learn more

* xref:{relative_path}storage_shelves_endpoint_overview.html[DOC
/storage/shelves]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|location_led
|string
|query
|False
a|Filter by location_led
```

```
* Introduced in: 9.10


|local
|boolean
|query
|False
a|Filter by local

* Introduced in: 9.8


|fans.rpm
|integer
|query
|False
a|Filter by fans.rpm

* Introduced in: 9.9


|fans.state
|string
|query
|False
a|Filter by fans.state

* Introduced in: 9.9


|fans.installed
|boolean
|query
|False
a|Filter by fans.installed

* Introduced in: 9.13


|fans.location
|string
|query
|False
a|Filter by fans.location

* Introduced in: 9.9
```

```
|fans.id
|integer
|query
|False
a|Filter by fans.id

* Introduced in: 9.9


|acps.subnet
|string
|query
|False
a|Filter by acps.subnet

* Introduced in: 9.10


|acps.error.type
|string
|query
|False
a|Filter by acps.error.type

* Introduced in: 9.10


|acps.error.reason.arguments.message
|string
|query
|False
a|Filter by acps.error.reason.arguments.message

* Introduced in: 9.10


|acps.error.reason.arguments.code
|string
|query
|False
a|Filter by acps.error.reason.arguments.code

* Introduced in: 9.10
```

```
|acps.error.reason.code
|string
|query
|False
a|Filter by acps.error.reason.code

* Introduced in: 9.10


|acps.error.reason.message
|string
|query
|False
a|Filter by acps.error.reason.message

* Introduced in: 9.10


|acps.error.severity
|string
|query
|False
a|Filter by acps.error.severity

* Introduced in: 9.10


|acps.address
|string
|query
|False
a|Filter by acps.address

* Introduced in: 9.10


|acps.netmask
|string
|query
|False
a|Filter by acps.netmask

* Introduced in: 9.10


|acps.node.name
|string
```

```
|query
|False
a|Filter by acps.node.name

* Introduced in: 9.10


|acps.node.uuid
|string
|query
|False
a|Filter by acps.node.uuid

* Introduced in: 9.10


|acps.enabled
|boolean
|query
|False
a|Filter by acps.enabled

* Introduced in: 9.10


|acps.connection_state
|string
|query
|False
a|Filter by acps.connection_state

* Introduced in: 9.10


|acps.port
|string
|query
|False
a|Filter by acps.port

* Introduced in: 9.10


|acps.channel
|string
|query
|False
```

```
a|Filter by acps.channel

* Introduced in: 9.10


|frus.firmware_version
|string
|query
|False
a|Filter by frus.firmware_version


|frus.type
|string
|query
|False
a|Filter by frus.type


|frus.id
|integer
|query
|False
a|Filter by frus.id


|frus.serial_number
|string
|query
|False
a|Filter by frus.serial_number


|frus.state
|string
|query
|False
a|Filter by frus.state


|frus.installed
|boolean
|query
|False
a|Filter by frus.installed

* Introduced in: 9.10
```

```
|frus.psu.crest_factor
|integer
|query
|False
a|Filter by frus.psu.crest_factor

* Introduced in: 9.10


|frus.psu.power_rating
|integer
|query
|False
a|Filter by frus.psu.power_rating

* Introduced in: 9.10


|frus.psu.power_drawn
|integer
|query
|False
a|Filter by frus.psu.power_drawn

* Introduced in: 9.10


|frus.psu.model
|string
|query
|False
a|Filter by frus.psu.model

* Introduced in: 9.10


|frus.part_number
|string
|query
|False
a|Filter by frus.part_number


|uid
|string
```

```
|query
|False
a|Filter by uid


|module_type
|string
|query
|False
a|Filter by module_type


|manufacturer.name
|string
|query
|False
a|Filter by manufacturer.name

* Introduced in: 9.10


|vendor.part_number
|string
|query
|False
a|Filter by vendor.part_number

* Introduced in: 9.8


|vendor.name
|string
|query
|False
a|Filter by vendor.name

* Introduced in: 9.10


|vendor.product
|string
|query
|False
a|Filter by vendor.product

* Introduced in: 9.8
```

```
|vendor.manufacturer
|string
|query
|False
a|Filter by vendor.manufacturer

* Introduced in: 9.8


|vendor.serial_number
|string
|query
|False
a|Filter by vendor.serial_number

* Introduced in: 9.8


|temperature_sensors.installed
|boolean
|query
|False
a|Filter by temperature_sensors.installed

* Introduced in: 9.13


|temperature_sensors.threshold.high.warning
|integer
|query
|False
a|Filter by temperature_sensors.threshold.high.warning

* Introduced in: 9.10


|temperature_sensors.threshold.high.critical
|integer
|query
|False
a|Filter by temperature_sensors.threshold.high.critical

* Introduced in: 9.10


|temperature_sensors.threshold.low.warning
```

```
|integer
|query
|False
a|Filter by temperature_sensors.threshold.low.warning

* Introduced in: 9.10


|temperature_sensors.threshold.low.critical
|integer
|query
|False
a|Filter by temperature_sensors.threshold.low.critical

* Introduced in: 9.10


|temperature_sensors.state
|string
|query
|False
a|Filter by temperature_sensors.state

* Introduced in: 9.10


|temperature_sensors.ambient
|boolean
|query
|False
a|Filter by temperature_sensors.ambient

* Introduced in: 9.10


|temperature_sensors.id
|integer
|query
|False
a|Filter by temperature_sensors.id

* Introduced in: 9.10


|temperature_sensors.location
|string
|query
```

```
|False
a|Filter by temperature_sensors.location

* Introduced in: 9.10


|temperature_sensors.temperature
|integer
|query
|False
a|Filter by temperature_sensors.temperature

* Introduced in: 9.10


|paths.node.name
|string
|query
|False
a|Filter by paths.node.name


|paths.node.uuid
|string
|query
|False
a|Filter by paths.node.uuid


|paths.name
|string
|query
|False
a|Filter by paths.name


|name
|string
|query
|False
a|Filter by name


|state
|string
|query
|False
```

```
a|Filter by state


|model
|string
|query
|False
a|Filter by model


|errors.reason.arguments.message
|string
|query
|False
a|Filter by errors.reason.arguments.message

* Introduced in: 9.10


|errors.reason.arguments.code
|string
|query
|False
a|Filter by errors.reason.arguments.code

* Introduced in: 9.10


|errors.reason.code
|string
|query
|False
a|Filter by errors.reason.code

* Introduced in: 9.9


|errors.reason.message
|string
|query
|False
a|Filter by errors.reason.message

* Introduced in: 9.9


|voltage_sensors.id
```

```
|integer
|query
|False
a|Filter by voltage_sensors.id

* Introduced in: 9.10


|voltage_sensors.location
|string
|query
|False
a|Filter by voltage_sensors.location

* Introduced in: 9.10


|voltage_sensors.voltage
|number
|query
|False
a|Filter by voltage_sensors.voltage

* Introduced in: 9.10


|voltage_sensors.installed
|boolean
|query
|False
a|Filter by voltage_sensors.installed

* Introduced in: 9.13


|voltage_sensors.state
|string
|query
|False
a|Filter by voltage_sensors.state

* Introduced in: 9.10


|bays.has_disk
|boolean
|query
```

```
|False
a|Filter by bays.has_disk


|bays.state
|string
|query
|False
a|Filter by bays.state


|bays.drawer.id
|integer
|query
|False
a|Filter by bays.drawer.id

* Introduced in: 9.11


|bays.drawer.slot
|integer
|query
|False
a|Filter by bays.drawer.slot

* Introduced in: 9.11


|bays.id
|integer
|query
|False
a|Filter by bays.id


|bays.type
|string
|query
|False
a|Filter by bays.type


|ports.internal
|boolean
|query
|False
```

```
a|Filter by ports.internal


|ports.module_id
|string
|query
|False
a|Filter by ports.module_id


|ports.id
|integer
|query
|False
a|Filter by ports.id


|ports.wwn
|string
|query
|False
a|Filter by ports.wwn


|ports.remote.chassis
|string
|query
|False
a|Filter by ports.remote.chassis


|ports.remote.device
|string
|query
|False
a|Filter by ports.remote.device

* Introduced in: 9.8


|ports.remote.wwn
|string
|query
|False
a|Filter by ports.remote.wwn
```

```
|ports.remote.port
|string
|query
|False
a|Filter by ports.remote.port


|ports.remote.phy
|string
|query
|False
a|Filter by ports.remote.phy


|ports.remote.mac_address
|string
|query
|False
a|Filter by ports.remote.mac_address


|ports.state
|string
|query
|False
a|Filter by ports.state


|ports.mac_address
|string
|query
|False
a|Filter by ports.mac_address


|ports.speed
|integer
|query
|False
a|Filter by ports.speed

* Introduced in: 9.17


|ports.cable.part_number
|string
|query
```

```
|False
a|Filter by ports.cable.part_number


|ports.cable.length
|string
|query
|False
a|Filter by ports.cable.length


|ports.cable.serial_number
|string
|query
|False
a|Filter by ports.cable.serial_number


|ports.cable.identifier
|string
|query
|False
a|Filter by ports.cable.identifier


|ports.designator
|string
|query
|False
a|Filter by ports.designator


|connection_type
|string
|query
|False
a|Filter by connection_type


|raw_capacity
|integer
|query
|False
a|Filter by raw_capacity

* Introduced in: 9.17
```

```
|drawers.part_number
|string
|query
|False
a|Filter by drawers.part_number


|drawers.error
|string
|query
|False
a|Filter by drawers.error


|drawers.state
|string
|query
|False
a|Filter by drawers.state


|drawers.closed
|boolean
|query
|False
a|Filter by drawers.closed


|drawers.serial_number
|string
|query
|False
a|Filter by drawers.serial_number


|drawers.id
|integer
|query
|False
a|Filter by drawers.id


|drawers.disk_count
|integer
|query
|False
```

```
a|Filter by drawers.disk_count


|id
|string
|query
|False
a|Filter by id


|disk_count
|integer
|query
|False
a|Filter by disk_count


|current_sensors.id
|integer
|query
|False
a|Filter by current_sensors.id

* Introduced in: 9.10


|current_sensors.current
|integer
|query
|False
a|Filter by current_sensors.current

* Introduced in: 9.10


|current_sensors.location
|string
|query
|False
a|Filter by current_sensors.location

* Introduced in: 9.10


|current_sensors.installed
|boolean
|query
```

```
|False
a|Filter by current_sensors.installed

* Introduced in: 9.13


|current_sensors.state
|string
|query
|False
a|Filter by current_sensors.state

* Introduced in: 9.10


|internal
|boolean
|query
|False
a|Filter by internal


|serial_number
|string
|query
|False
a|Filter by serial_number


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
```

```
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records
```

```
|records
|array[link:#shelf[shelf]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "acps": [
        {
          "address": "192.168.1.104",
          "channel": "out_of_band",
          "connection_state": "full_connectivity",
          "error": {
            "reason": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
              ],
              "code": "4",
              "message": "entry doesn't exist"
            },
            "severity": "string",
            "type": "string"
          },
          "netmask": "255.255.252.0",
          "node": {
```

```
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "port": "e0P",
          "subnet": "192.168.0.1"
        }
      ],
      "bays": [
        {
          "drawer": {
            "id": 1,
            "slot": 0
          },
          "id": 0,
          "state": "ok",
          "type": "single_disk"
        }
      ],
      "connection_type": "sas",
      "current_sensors": [
        {
          "current": 14410,
          "id": 1,
          "installed": 1,
          "location": "rear of the shelf on the lower left power supply",
          "state": "ok"
        }
      ],
      "disk_count": 12,
      "drawers": [
        {
          "disk_count": 12,
          "error": "string",
          "part_number": "111-03071",
          "serial_number": "021604008263",
          "state": "ok"
        }
      ],
      "errors": [
        {
          "reason": {
```

```
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          }
        }
      ],
      "fans": [
        {
          "id": 1,
          "installed": 1,
          "location": "rear of the shelf on the lower left power supply",
          "rpm": 3020,
          "state": "ok"
        }
      ],
      "frus": [
        {
          "firmware_version": "0191",
          "installed": 1,
          "part_number": "111-00690+A2",
          "psu": {
            "crest_factor": 92,
            "model": "00",
            "power_drawn": 210,
            "power_rating": 1600
          },
          "serial_number": "8000166294",
          "state": "error",
          "type": "module"
        }
      ],
      "id": "1",
      "location_led": "off",
      "manufacturer": {
        "name": "NETAPP"
      },
      "model": "DS2246",
      "module_type": "iom6",
      "name": "1.1",
      "paths": [
        {
```

```
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "2a",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "ports": [
        {
          "cable": {
            "identifier": "500a0980000b6c3f-50000d1703544b80",
            "length": "2m",
            "part_number": "112-00431+A0",
            "serial_number": "616930439"
          },
          "designator": "square",
          "id": 0,
          "mac_address": "string",
          "module_id": "a",
          "remote": {
            "chassis": "string",
            "device": "string",
            "mac_address": "string",
            "phy": "12",
            "port": "string",
            "wwn": "50000D1703544B80"
          },
          "speed": 100,
          "state": "connected",
          "wwn": "500A0980000B6C3F"
        }
      ],
      "raw_capacity": 5157308772,
      "serial_number": "SHFMS1514000895",
      "state": "ok",
      "temperature_sensors": [
```

```
        {
          "ambient": "",
          "id": 1,
          "installed": 1,
          "location": "temp sensor on midplane left",
          "state": "ok",
          "temperature": 32,
          "threshold": {
            "high": {
              "critical": 60,
              "warning": 55
            },
            "low": {
              "critical": 0,
              "warning": 5
            }
          }
        }
      ],
      "uid": "7777841915827391056",
      "vendor": {
        "manufacturer": "XYZ",
        "name": "XYZ",
        "part_number": "A92831142733",
        "product": "LS2246",
        "serial_number": "891234572210221"
      },
      "voltage_sensors": [
        {
          "id": 1,
          "installed": 1,
          "location": "rear of the shelf on the lower left power supply",
          "state": "ok",
          "voltage": 12.18
        }
      ]
    }
  ]
}
====

== Error
```


Status: Default, Error

```
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
```

```
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#error]
[.api-collapsible-fifth-title]
error

Error object is populated when connection_state becomes non-optimal


[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|reason
|link:#error[error]
a|

|severity
|string
a|

|type
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|

|uuid
|string
a|

|===


[#acps]
[.api-collapsible-fifth-title]
acps

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|

|channel
|string
a|

|connection_state
|string
a|

|enabled
|boolean
a|

|error
|link:#error[error]
a|Error object is populated when connection_state becomes non-optimal


|netmask
|string
a|

|node
```

```
|link:#node[node]
a|

|port
|string
a|

|subnet
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|The drawer containing this bay


|slot
|integer
a|The drawer slot for this bay


|===


[#bays]
[.api-collapsible-fifth-title]
bays

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|drawer
|link:#drawer[drawer]
a|

|has_disk
|boolean
a|

|id
|integer
a|

|state
|string
a|

|type
|string
a|

|===


[#current_sensors]
[.api-collapsible-fifth-title]
current_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|current
|integer
a|Current, in milliamps


|id
|integer
a|

|installed
|boolean
a|

|location
```

```
|string
a|

|state
|string
a|

|===


[#drawers]
[.api-collapsible-fifth-title]
drawers

[cols=3*,options=header]
|===
|Name
|Type
|Description

|closed
|boolean
a|

|disk_count
|integer
a|

|error
|string
a|

|id
|integer
a|

|part_number
|string
a|

|serial_number
|string
a|

|state
|string
a|
```

```
|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|

|===


[#fans]
[.api-collapsible-fifth-title]
fans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|rpm
|integer
a|
```

```
|state
|string
a|

|===


[#psu]
[.api-collapsible-fifth-title]
psu

[cols=3*,options=header]
|===
|Name
|Type
|Description

|crest_factor
|integer
a|The ratio of the peak voltage to the root-mean-square voltage


|model
|string
a|

|power_drawn
|integer
a|Power drawn, in watts


|power_rating
|integer
a|Power rating, in watts


|===


[#frus]
[.api-collapsible-fifth-title]
frus

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|firmware_version
|string
a|

|id
|integer
a|

|installed
|boolean
a|

|part_number
|string
a|

|psu
|link:#psu[psu]
a|

|serial_number
|string
a|

|state
|string
a|

|type
|string
a|

|===


[#manufacturer]
[.api-collapsible-fifth-title]
manufacturer

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|

|===


[#paths]
[.api-collapsible-fifth-title]
paths

Storage port


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|node
|link:#node[node]
a|

|===


[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
```

```
a|

|length
|string
a|

|part_number
|string
a|

|serial_number
|string
a|

|===


[#remote]
[.api-collapsible-fifth-title]
remote

[cols=3*,options=header]
|===
|Name
|Type
|Description

|chassis
|string
a|

|device
|string
a|

|mac_address
|string
a|

|phy
|string
a|

|port
|string
a|
```

```
|wwn
|string
a|

|===


[#ports]
[.api-collapsible-fifth-title]
ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cable
|link:#cable[cable]
a|

|designator
|string
a|

|id
|integer
a|

|internal
|boolean
a|

|mac_address
|string
a|

|module_id
|string
a|

|remote
|link:#remote[remote]
a|

|speed
|integer
```

```
a|

|state
|string
a|

|wwn
|string
a|

|===


[#high]
[.api-collapsible-fifth-title]
high

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
|integer
a|High critical threshold, in degrees Celsius


|warning
|integer
a|High warning threshold, in degrees Celsius


|===


[#low]
[.api-collapsible-fifth-title]
low

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
```

```
|integer
a|Low critical threshold, in degrees Celsius


|warning
|integer
a|Low warning threshold, in degrees Celsius


|===


[#threshold]
[.api-collapsible-fifth-title]
threshold

[cols=3*,options=header]
|===
|Name
|Type
|Description

|high
|link:#high[high]
a|

|low
|link:#low[low]
a|

|===


[#temperature_sensors]
[.api-collapsible-fifth-title]
temperature_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ambient
|boolean
a|Sensor that measures the ambient temperature
```

```
|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|temperature
|integer
a|Temperature, in degrees Celsius


|threshold
|link:#threshold[threshold]
a|

|===


[#vendor]
[.api-collapsible-fifth-title]
vendor

[cols=3*,options=header]
|===
|Name
|Type
|Description

|manufacturer
|string
a|Support for this field will be removed in a future release. Please use
vendor.name for this field.


|name
|string
```

```
a|

|part_number
|string
a|Part number


|product
|string
a|Product name


|serial_number
|string
a|Serial number


|===


[#voltage_sensors]
[.api-collapsible-fifth-title]
voltage_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|voltage
```

```
|number
a|Voltage, in volts


|===


[#shelf]
[.api-collapsible-fifth-title]
shelf

[cols=3*,options=header]
|===
|Name
|Type
|Description

|acps
|array[link:#acps[acps]]
a|Alternate Control Paths to ACP processors/functions in shelf modules and
expanders


|bays
|array[link:#bays[bays]]
a|

|connection_type
|string
a|

|current_sensors
|array[link:#current_sensors[current_sensors]]
a|

|disk_count
|integer
a|

|drawers
|array[link:#drawers[drawers]]
a|

|errors
|array[link:#errors[errors]]
a|
```

```
|fans
|array[link:#fans[fans]]
a|

|frus
|array[link:#frus[frus]]
a|

|id
|string
a|

|internal
|boolean
a|

|local
|boolean
a|

|location_led
|string
a|

|manufacturer
|link:#manufacturer[manufacturer]
a|

|model
|string
a|

|module_type
|string
a|

|name
|string
a|

|paths
|array[link:#paths[paths]]
a|

|ports
|array[link:#ports[ports]]
a|
```

```
|raw_capacity
|integer
a|Raw capacity of drives in shelf in 4k blocks


|serial_number
|string
a|

|state
|string
a|

|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|

|uid
|string
a|

|vendor
|link:#vendor[vendor]
a|

|voltage_sensors
|array[link:#voltage_sensors[voltage_sensors]]
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDb91a95951308429d051007a0911f3dd7]]
= Retrieve a shelf

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/shelves/{uid}`#

*Introduced In:* 9.6

Retrieves a specific shelf.

== Related ONTAP commands

* `storage shelf show`
* `storage shelf port show`
* `storage shelf drawer show`
* `storage shelf drawer show-slot`
* `storage shelf acp show`

== Learn more

* xref:{relative_path}storage_shelves_endpoint_overview.html[DOC
/storage/shelves]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uid
|string
|path
|True
a|Shelf UID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|acps
|array[link:#acps[acps]]
a|Alternate Control Paths to ACP processors/functions in shelf modules and
expanders


|bays
|array[link:#bays[bays]]
a|
```

```
|connection_type
|string
a|

|current_sensors
|array[link:#current_sensors[current_sensors]]
a|

|disk_count
|integer
a|

|drawers
|array[link:#drawers[drawers]]
a|

|errors
|array[link:#errors[errors]]
a|

|fans
|array[link:#fans[fans]]
a|

|frus
|array[link:#frus[frus]]
a|

|id
|string
a|

|internal
|boolean
a|

|local
|boolean
a|

|location_led
|string
a|

|manufacturer
|link:#manufacturer[manufacturer]
```

```
a|

|model
|string
a|

|module_type
|string
a|

|name
|string
a|

|paths
|array[link:#paths[paths]]
a|

|ports
|array[link:#ports[ports]]
a|

|raw_capacity
|integer
a|Raw capacity of drives in shelf in 4k blocks


|serial_number
|string
a|

|state
|string
a|

|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|

|uid
|string
a|

|vendor
|link:#vendor[vendor]
a|
```

```
|voltage_sensors
|array[link:#voltage_sensors[voltage_sensors]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "acps": [
    {
      "address": "192.168.1.104",
      "channel": "out_of_band",
      "connection_state": "full_connectivity",
      "error": {
        "reason": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "severity": "string",
        "type": "string"
      },
      "netmask": "255.255.252.0",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": "e0P",
      "subnet": "192.168.0.1"
    }
  ],
  "bays": [
```

```
    {
      "drawer": {
        "id": 1,
        "slot": 0
      },
      "id": 0,
      "state": "ok",
      "type": "single_disk"
    }
  ],
  "connection_type": "sas",
  "current_sensors": [
    {
      "current": 14410,
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "state": "ok"
    }
  ],
  "disk_count": 12,
  "drawers": [
    {
      "disk_count": 12,
      "error": "string",
      "part_number": "111-03071",
      "serial_number": "021604008263",
      "state": "ok"
    }
  ],
  "errors": [
    {
      "reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    }
  ],
  "fans": [
    {
```

```
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "rpm": 3020,
      "state": "ok"
    }
  ],
  "frus": [
    {
      "firmware_version": "0191",
      "installed": 1,
      "part_number": "111-00690+A2",
      "psu": {
        "crest_factor": 92,
        "model": "00",
        "power_drawn": 210,
        "power_rating": 1600
      },
      "serial_number": "8000166294",
      "state": "error",
      "type": "module"
    }
  ],
  "id": "1",
  "location_led": "off",
  "manufacturer": {
    "name": "NETAPP"
  },
  "model": "DS2246",
  "module_type": "iom6",
  "name": "1.1",
  "paths": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "2a",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
```

```
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "ports": [
    {
      "cable": {
        "identifier": "500a0980000b6c3f-50000d1703544b80",
        "length": "2m",
        "part_number": "112-00431+A0",
        "serial_number": "616930439"
      },
      "designator": "square",
      "id": 0,
      "mac_address": "string",
      "module_id": "a",
      "remote": {
        "chassis": "string",
        "device": "string",
        "mac_address": "string",
        "phy": "12",
        "port": "string",
        "wwn": "50000D1703544B80"
      },
      "speed": 100,
      "state": "connected",
      "wwn": "500A0980000B6C3F"
    }
  ],
  "raw_capacity": 5157308772,
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "temperature_sensors": [
    {
      "ambient": "",
      "id": 1,
      "installed": 1,
      "location": "temp sensor on midplane left",
      "state": "ok",
      "temperature": 32,
      "threshold": {
        "high": {
          "critical": 60,
          "warning": 55
        },
        "low": {
```

```
          "critical": 0,
          "warning": 5
        }
      }
    }
  ],
  "uid": "7777841915827391056",
  "vendor": {
    "manufacturer": "XYZ",
    "name": "XYZ",
    "part_number": "A92831142733",
    "product": "LS2246",
    "serial_number": "891234572210221"
  },
  "voltage_sensors": [
    {
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "state": "ok",
      "voltage": 12.18
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
```

```
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#error]
[.api-collapsible-fifth-title]
error

Error object is populated when connection_state becomes non-optimal


[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|

|severity
|string
```

```
a|

|type
|string
a|

|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#acps]
[.api-collapsible-fifth-title]
acps

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|

|channel
|string
a|

|connection_state
|string
a|

|enabled
|boolean
a|
```

```
|error
|link:#error[error]
a|Error object is populated when connection_state becomes non-optimal


|netmask
|string
a|

|node
|link:#node[node]
a|

|port
|string
a|

|subnet
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|The drawer containing this bay


|slot
|integer
a|The drawer slot for this bay


|===
```

```
[#bays]
[.api-collapsible-fifth-title]
bays

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drawer
|link:#drawer[drawer]
a|

|has_disk
|boolean
a|

|id
|integer
a|

|state
|string
a|

|type
|string
a|

|===


[#current_sensors]
[.api-collapsible-fifth-title]
current_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|current
|integer
a|Current, in milliamps
```

```
|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|===


[#drawers]
[.api-collapsible-fifth-title]
drawers

[cols=3*,options=header]
|===
|Name
|Type
|Description

|closed
|boolean
a|

|disk_count
|integer
a|

|error
|string
a|

|id
|integer
a|

|part_number
```

```
|string
a|

|serial_number
|string
a|

|state
|string
a|

|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|reason
|link:#error[error]
a|

|===


[#fans]
[.api-collapsible-fifth-title]
fans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
```

```
a|

|location
|string
a|

|rpm
|integer
a|

|state
|string
a|

|===


[#psu]
[.api-collapsible-fifth-title]
psu

[cols=3*,options=header]
|===
|Name
|Type
|Description

|crest_factor
|integer
a|The ratio of the peak voltage to the root-mean-square voltage


|model
|string
a|

|power_drawn
|integer
a|Power drawn, in watts


|power_rating
|integer
a|Power rating, in watts


|===
```

```
[#frus]
[.api-collapsible-fifth-title]
frus

[cols=3*,options=header]
|===
|Name
|Type
|Description

|firmware_version
|string
a|

|id
|integer
a|

|installed
|boolean
a|

|part_number
|string
a|

|psu
|link:#psu[psu]
a|

|serial_number
|string
a|

|state
|string
a|

|type
|string
a|

|===
```

```
[#manufacturer]
[.api-collapsible-fifth-title]
manufacturer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|===


[#paths]
[.api-collapsible-fifth-title]
paths

Storage port


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|node
|link:#node[node]
a|

|===


[#cable]
[.api-collapsible-fifth-title]
```

```
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
a|

|length
|string
a|

|part_number
|string
a|

|serial_number
|string
a|

|===


[#remote]
[.api-collapsible-fifth-title]
remote

[cols=3*,options=header]
|===
|Name
|Type
|Description

|chassis
|string
a|

|device
|string
a|

|mac_address
|string
```

```
a|

|phy
|string
a|

|port
|string
a|

|wwn
|string
a|

|===


[#ports]
[.api-collapsible-fifth-title]
ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cable
|link:#cable[cable]
a|

|designator
|string
a|

|id
|integer
a|

|internal
|boolean
a|

|mac_address
|string
a|
```

```
|module_id
|string
a|

|remote
|link:#remote[remote]
a|

|speed
|integer
a|

|state
|string
a|

|wwn
|string
a|

|===


[#high]
[.api-collapsible-fifth-title]
high

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
|integer
a|High critical threshold, in degrees Celsius


|warning
|integer
a|High warning threshold, in degrees Celsius


|===


[#low]
```

```
[.api-collapsible-fifth-title]
low

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
|integer
a|Low critical threshold, in degrees Celsius


|warning
|integer
a|Low warning threshold, in degrees Celsius


|===


[#threshold]
[.api-collapsible-fifth-title]
threshold

[cols=3*,options=header]
|===
|Name
|Type
|Description

|high
|link:#high[high]
a|

|low
|link:#low[low]
a|

|===


[#temperature_sensors]
[.api-collapsible-fifth-title]
temperature_sensors
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|ambient
|boolean
a|Sensor that measures the ambient temperature


|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|temperature
|integer
a|Temperature, in degrees Celsius


|threshold
|link:#threshold[threshold]
a|

|===


[#vendor]
[.api-collapsible-fifth-title]
vendor

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|manufacturer
|string
a|Support for this field will be removed in a future release. Please use
vendor.name for this field.


|name
|string
a|

|part_number
|string
a|Part number


|product
|string
a|Product name


|serial_number
|string
a|Serial number


|===


[#voltage_sensors]
[.api-collapsible-fifth-title]
voltage_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
a|
```

```
|location
|string
a|

|state
|string
a|

|voltage
|number
a|Voltage, in volts


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID4161c9438f750cf5fd14504662b1c705]]
= Update a shelf location LED

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/shelves/{uid}`#

*Introduced In:* 9.10

Updates a shelf location LED.

== Related ONTAP commands

* `storage shelf location-led modify`

== Learn more

* xref:{relative_path}storage_shelves_endpoint_overview.html[DOC
/storage/shelves]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uid
|string
|path
|True
a|Shelf UID

|===
```

```
== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|acps
|array[link:#acps[acps]]
a|Alternate Control Paths to ACP processors/functions in shelf modules and
expanders


|connection_type
|string
a|

|current_sensors
|array[link:#current_sensors[current_sensors]]
a|

|disk_count
|integer
a|

|drawers
|array[link:#drawers[drawers]]
a|

|fans
|array[link:#fans[fans]]
a|

|frus
|array[link:#frus[frus]]
a|

|id
|string
a|

|internal
|boolean
a|
```

```
|local
|boolean
a|

|location_led
|string
a|

|manufacturer
|link:#manufacturer[manufacturer]
a|

|model
|string
a|

|module_type
|string
a|

|name
|string
a|

|paths
|array[link:#paths[paths]]
a|

|raw_capacity
|integer
a|Raw capacity of drives in shelf in 4k blocks


|serial_number
|string
a|

|state
|string
a|

|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|

|uid
```

```
|string
a|

|vendor
|link:#vendor[vendor]
a|

|voltage_sensors
|array[link:#voltage_sensors[voltage_sensors]]
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "acps": [
    {
      "address": "192.168.1.104",
      "channel": "out_of_band",
      "connection_state": "full_connectivity",
      "error": {
        "severity": "string",
        "type": "string"
      },
      "netmask": "255.255.252.0",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": "e0P",
      "subnet": "192.168.0.1"
    }
  ],
  "connection_type": "sas",
  "current_sensors": [
    {
      "current": 14410,
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "state": "ok"
    }
```

```
  ],
  "disk_count": 12,
  "drawers": [
    {
      "disk_count": 12,
      "error": "string",
      "part_number": "111-03071",
      "serial_number": "021604008263",
      "state": "ok"
    }
  ],
  "fans": [
    {
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "rpm": 3020,
      "state": "ok"
    }
  ],
  "frus": [
    {
      "firmware_version": "0191",
      "installed": 1,
      "part_number": "111-00690+A2",
      "psu": {
        "crest_factor": 92,
        "model": "00",
        "power_drawn": 210,
        "power_rating": 1600
      },
      "serial_number": "8000166294",
      "state": "error",
      "type": "module"
    }
  ],
  "id": "1",
  "location_led": "off",
  "manufacturer": {
    "name": "NETAPP"
  },
  "model": "DS2246",
  "module_type": "iom6",
  "name": "1.1",
  "paths": [
    {
```

```
      "name": "2a",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "raw_capacity": 5157308772,
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "temperature_sensors": [
    {
      "ambient": "",
      "id": 1,
      "installed": 1,
      "location": "temp sensor on midplane left",
      "state": "ok",
      "temperature": 32,
      "threshold": {
        "high": {
          "critical": 60,
          "warning": 55
        },
        "low": {
          "critical": 0,
          "warning": 5
        }
      }
    }
  ],
  "uid": "7777841915827391056",
  "vendor": {
    "manufacturer": "XYZ",
    "name": "XYZ",
    "part_number": "A92831142733",
    "product": "LS2246",
    "serial_number": "891234572210221"
  },
  "voltage_sensors": [
    {
      "id": 1,
      "installed": 1,
      "location": "rear of the shelf on the lower left power supply",
      "state": "ok",
      "voltage": 12.18
    }
```

```
  ]
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 17825872
| Shelf locate request failed because shelf "+++<name>+++\" was not
found.+++</name>+++

| 17825873
| Shelf locate request failed because shelf "+++<name>+++\" does not
support this command.+++</name>+++

| 17825874
| Shelf locate request failed for shelf "+++<name>+++\" with an unknown
error.+++</name>+++

| 17825875
| Shelf locate request failed for shelf "+++<name>+++\" because shelf
modules are unreachable.+++</name>+++
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|===


[#error]
[.api-collapsible-fifth-title]
error

Error object is populated when connection_state becomes non-optimal


[cols=3*,options=header]
|===
|Name
|Type
|Description

|severity
|string
a|

|type
|string
a|

|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
```

```
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#acps]
[.api-collapsible-fifth-title]
acps

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|

|channel
|string
a|

|connection_state
|string
a|

|enabled
|boolean
```

```
a|

|error
|link:#error[error]
a|Error object is populated when connection_state becomes non-optimal


|netmask
|string
a|

|node
|link:#node[node]
a|

|port
|string
a|

|subnet
|string
a|

|===


[#drawer]
[.api-collapsible-fifth-title]
drawer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|The drawer containing this bay


|slot
|integer
a|The drawer slot for this bay


|===
```

```
[#bays]
[.api-collapsible-fifth-title]
bays

[cols=3*,options=header]
|===
|Name
|Type
|Description

|drawer
|link:#drawer[drawer]
a|

|has_disk
|boolean
a|

|id
|integer
a|

|state
|string
a|

|type
|string
a|

|===


[#current_sensors]
[.api-collapsible-fifth-title]
current_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|current
|integer
```

```
a|Current, in milliamps


|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|===


[#drawers]
[.api-collapsible-fifth-title]
drawers

[cols=3*,options=header]
|===
|Name
|Type
|Description

|closed
|boolean
a|

|disk_count
|integer
a|

|error
|string
a|

|id
|integer
a|
```

```
|part_number
|string
a|

|serial_number
|string
a|

|state
|string
a|

|===


[#errors]
[.api-collapsible-fifth-title]
errors
[#fans]
[.api-collapsible-fifth-title]
fans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|rpm
|integer
a|

|state
|string
```

```
a|

|===


[#psu]
[.api-collapsible-fifth-title]
psu

[cols=3*,options=header]
|===
|Name
|Type
|Description

|crest_factor
|integer
a|The ratio of the peak voltage to the root-mean-square voltage


|model
|string
a|

|power_drawn
|integer
a|Power drawn, in watts


|power_rating
|integer
a|Power rating, in watts


|===


[#frus]
[.api-collapsible-fifth-title]
frus

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|firmware_version
|string
a|

|id
|integer
a|

|installed
|boolean
a|

|part_number
|string
a|

|psu
|link:#psu[psu]
a|

|serial_number
|string
a|

|state
|string
a|

|type
|string
a|

|===


[#manufacturer]
[.api-collapsible-fifth-title]
manufacturer

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
```

```
a|

|===


[#paths]
[.api-collapsible-fifth-title]
paths

Storage port


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|===


[#cable]
[.api-collapsible-fifth-title]
cable

[cols=3*,options=header]
|===
|Name
|Type
|Description

|identifier
|string
a|

|length
|string
a|
```

```
|part_number
|string
a|

|serial_number
|string
a|

|===


[#remote]
[.api-collapsible-fifth-title]
remote

[cols=3*,options=header]
|===
|Name
|Type
|Description

|chassis
|string
a|

|device
|string
a|

|mac_address
|string
a|

|phy
|string
a|

|port
|string
a|

|wwn
|string
a|

|===
```

```
[#ports]
[.api-collapsible-fifth-title]
ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cable
|link:#cable[cable]
a|

|designator
|string
a|

|id
|integer
a|

|internal
|boolean
a|

|mac_address
|string
a|

|module_id
|string
a|

|remote
|link:#remote[remote]
a|

|speed
|integer
a|

|state
|string
a|
```

```
|wwn
|string
a|

|===


[#high]
[.api-collapsible-fifth-title]
high

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
|integer
a|High critical threshold, in degrees Celsius


|warning
|integer
a|High warning threshold, in degrees Celsius


|===


[#low]
[.api-collapsible-fifth-title]
low

[cols=3*,options=header]
|===
|Name
|Type
|Description

|critical
|integer
a|Low critical threshold, in degrees Celsius


|warning
|integer
```

```
a|Low warning threshold, in degrees Celsius


|===


[#threshold]
[.api-collapsible-fifth-title]
threshold

[cols=3*,options=header]
|===
|Name
|Type
|Description

|high
|link:#high[high]
a|

|low
|link:#low[low]
a|

|===


[#temperature_sensors]
[.api-collapsible-fifth-title]
temperature_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ambient
|boolean
a|Sensor that measures the ambient temperature


|id
|integer
a|

|installed
```

```
|boolean
a|

|location
|string
a|

|state
|string
a|

|temperature
|integer
a|Temperature, in degrees Celsius


|threshold
|link:#threshold[threshold]
a|

|===


[#vendor]
[.api-collapsible-fifth-title]
vendor

[cols=3*,options=header]
|===
|Name
|Type
|Description

|manufacturer
|string
a|Support for this field will be removed in a future release. Please use
vendor.name for this field.


|name
|string
a|

|part_number
|string
a|Part number
```

```
|product
|string
a|Product name


|serial_number
|string
a|Serial number


|===


[#voltage_sensors]
[.api-collapsible-fifth-title]
voltage_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|

|installed
|boolean
a|

|location
|string
a|

|state
|string
a|

|voltage
|number
a|Voltage, in volts


|===
```

```
[#shelf]
[.api-collapsible-fifth-title]
shelf

[cols=3*,options=header]
|===
|Name
|Type
|Description

|acps
|array[link:#acps[acps]]
a|Alternate Control Paths to ACP processors/functions in shelf modules and
expanders


|connection_type
|string
a|

|current_sensors
|array[link:#current_sensors[current_sensors]]
a|

|disk_count
|integer
a|

|drawers
|array[link:#drawers[drawers]]
a|

|fans
|array[link:#fans[fans]]
a|

|frus
|array[link:#frus[frus]]
a|

|id
|string
a|

|internal
|boolean
```

```
a|

|local
|boolean
a|

|location_led
|string
a|

|manufacturer
|link:#manufacturer[manufacturer]
a|

|model
|string
a|

|module_type
|string
a|

|name
|string
a|

|paths
|array[link:#paths[paths]]
a|

|raw_capacity
|integer
a|Raw capacity of drives in shelf in 4k blocks


|serial_number
|string
a|

|state
|string
a|

|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|
```

```
|uid
|string
a|

|vendor
|link:#vendor[vendor]
a|

|voltage_sensors
|array[link:#voltage_sensors[voltage_sensors]]
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===
```

```
//end collapsible .Definitions block
====



:leveloffset: -1


= Manage snapshot policies

:leveloffset: +1


[[ID31efa2038c2e675680c4a57c5a876d93]]
= Manage snapshot policies



== Overview

In ONTAP, scheduled snapshot creation works based on snapshot policies.
ONTAP provides three cluster-wide snapshot policies: "default", "default-
1weekly" and "none".
A snapshot policy can have more than one schedule associated with it.
A snapshot policy can be linked to a storage object and based on the
schedule in the policy, snapshots will be created on the object at that
interval.
Each schedule in a snapshot policy has a snapshot name prefix attached to
it. Every snapshot created using this policy will have this prefix in its
name.
There is also a retention count associated with every schedule. This count
indicates the maximum number of snapshots that can exist for a given
schedule. Once the snapshot count reaches the retention count, on the next
create operation, the oldest snapshot is deleted.
A retention period can be associated with every schedule. During snapshot
creation, this period is set as SnapLock expiry time on snapshot locking
enabled volumes.

== Snapshot policy APIs

The following APIs are used to perform operations related to snapshot
policy information:

&ndash; POST      /api/storage/snapshot-policies

&ndash; GET       /api/storage/snapshot-policies
```

```
&ndash; GET       /api/storage/snapshot-policies/\{uuid}

&ndash; PATCH     /api/storage/snapshot-policies/\{uuid}

&ndash; DELETE    /api/storage/snapshot-policies/\{uuid}

== Examples

=== Creating a snapshot policy

The POST operation is used to create a snapshot policy with the specified
attributes.

----

# The API:
/api/storage/snapshot-policies

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/snapshot-policies" -H
"accept: application/hal+json" -d '{"name": "new_policy", "enabled":
"true", "comment": "policy comment", "copies": [{ "schedule": { "name":
"5min" }, "count": "5", "prefix": "xyz" , "retention_period": "PT20M" }],
"svm": { "name": "vs0"}}'

# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
    "svm": {
      "name": "vs0"
    },
    "name": "new_policy",
    "comment": "This is a 5min schedule policy",
```

```
    "enabled": true,
    "copies": [
      {
        "count": 5,
        "snapmirror_label": "-",
        "retention_period": "PT20M",
        "schedule": {
          "name": "5min"
        }
      }
    ]
  }
]
}
----

=== Retrieving snapshot policy attributes

The GET operation is used to retrieve snapshot policy attributes.

----

# The API:
/api/storage/snapshot-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/" -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:17:17 GMT
Server: libzapid-http
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 686
Content-Type: application/json
{
"records": [
  {
    "uuid": "0fa7a554-348d-11e9-b55e-005056bbf1c8",
    "name": "spsv0",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot-policies/0fa7a554-348d-11e9-b55e-
005056bbf1c8"
```

```
      }
    }
  },
  {
    "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c1c1656-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default-1weekly",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot-policies/3c1c1656-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
    "name": "none",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot-policies/3c228b82-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/snapshot-policies/"
  }
}
}
----

=== Retrieving the attributes of a specific snapshot policy

The GET operation is used to retrieve the attributes of a specific
```

```
snapshot policy.

----

# The API:
/api/storage/snapshot-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/3c112527-
2fe8-11e9-b55e-005056bbf1c8" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:24:48 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 381
Content-Type: application/json
{
"uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
"name": "default",
"comment": "Default policy with hourly, daily & weekly schedules.",
"enabled": true,
"scope": "cluster",
"copies": [
  {
    "count": 6,
    "prefix": "hourly",
    "schedule": {
      "name": "hourly"
    }
  },
  {
    "count": 2,
    "prefix": "daily",
    "schedule": {
      "name": "daily"
    }
  },
  {
    "count": 2,
    "prefix": "weekly",
    "schedule": {
      "name": "weekly"
    }
```

```
  }
],
"_links": {
  "self": {
    "href": "/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-
005056bbf1c8"
  }
}
}
----

=== Updating a snapshot policy

The PATCH operation is used to update the specific attributes of a
snapshot policy.

----

# The API:
/api/storage/snapshot-policies/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/snapshot-policies/ae9e65c4-
4506-11e9-aa44-005056bbc848" -d '{"enabled": "false" }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:27:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
----

=== Deleting a snapshot policy

The DELETE operation is used to delete a snapshot policy.

----

# The API:
/api/storage/snapshot-policies/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot-policies/ae9e65c4-
```

```
4506-11e9-aa44-005056bbc848" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:19:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
----





[[IDb484aaee82d5c367c0569b89b44b45f9]]
= Retrieve snapshot policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snapshot-policies`#

*Introduced In:* 9.6

Retrieves a collection of snapshot policies.

== Related ONTAP commands

* `snapshot policy show`

== Learn more

* xref:{relative_path}storage_snapshot-policies_endpoint_overview.html[DOC
/storage/snapshot-policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled


|name
|string
|query
|False
a|Filter by name


|svm.name
|string
|query
|False
a|Filter by svm.name


|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|uuid
|string
|query
|False
a|Filter by uuid


|copies.count
|integer
|query
|False
a|Filter by copies.count


|copies.snapmirror_label
|string
|query
|False
a|Filter by copies.snapmirror_label
```

```
|copies.schedule.name
|string
|query
|False
a|Filter by copies.schedule.name


|copies.schedule.uuid
|string
|query
|False
a|Filter by copies.schedule.uuid

* Introduced in: 9.14


|copies.retention_period
|string
|query
|False
a|Filter by copies.retention_period

* Introduced in: 9.12


|copies.prefix
|string
|query
|False
a|Filter by copies.prefix


|scope
|string
|query
|False
a|Filter by scope


|comment
|string
|query
|False
a|Filter by comment
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snapshot_policy[snapshot_policy]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
```

```
        }
      },
      "comment": "string",
      "copies": [
        {
          "prefix": "string",
          "retention_period": "string",
          "schedule": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "snapmirror_label": "string"
        }
      ],
      "enabled": 1,
      "name": "default",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#schedule]
[.api-collapsible-fifth-title]
schedule
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#copies]
[.api-collapsible-fifth-title]
copies

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule. The retention
```

```
period value represents a duration and must be specified in the ISO-8601
duration format. The retention period can be in years, months, days,
hours, and minutes. A period specified for years, months, and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively, for example "P10Y" represents a duration of
10 years. A duration in hours and minutes is represented by
"PT+++<num>+++H" and "PT+++<num>+++M" respectively. The period string must
contain only a single time element that is, either years, months, days,
hours, or minutes. A duration which combines different periods is not
supported, for example "P1Y10M" is not
supported.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

The snapshot policy object is associated with a read-write volume used to
create and delete snapshots at regular intervals.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A comment associated with the snapshot policy.


|copies
|array[link:#copies[copies]]
a|

|enabled
|boolean
a|Is the snapshot policy enabled?


|name
|string
a|Name of the snapshot policy.


|scope
```

```
|string
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID807269281da6abe3adf08aaf463be271]]
= Create a snapshot policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snapshot-policies`#

*Introduced In:* 9.6

Creates a snapshot policy.

== Required properties

* `svm.uuid` or `svm.name` - Specifies an SVM for policy creation. If not
specified, the snapshot policy will be created on the cluster admin SVM.
* `name` - Name for the snapshot policy.
* `copies.schedule` - Schedule name at which snapshots are captured on the
```

```
volume.
* `copies.count` - Number of snapshots to maintain for this schedule.

== Recommended optional properties

* `copies.prefix` - Prefix to use when creating snapshots at regular
intervals.
* `copies.snapmirror_label` - Label for SnapMirror operations.
* `copies.retention_period` - Retention period for snapshot locking
enabled volumes.The duration must be specified in ISO format or
"infinite".

== Default property values

If not specified in POST, the following default property values are
assigned:

* `svm.uuid` or `svm.name` - If not specified, the snapshot policy will be
created on the cluster admin SVM.
* `enabled` - _true_
* `copies.prefix` - Value of `schedule.name`

== Related ONTAP commands

* `snapshot policy create`

== Learn more

* xref:{relative_path}storage_snapshot-policies_endpoint_overview.html[DOC
/storage/snapshot-policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
```

```
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|A comment associated with the snapshot policy.


|copies
|array[link:#copies[copies]]
a|

|enabled
|boolean
a|Is the snapshot policy enabled?


|name
|string
a|Name of the snapshot policy.


|scope
|string
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
```

```
|string
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "retention_period": "string",
      "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====
```

== Response

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1638407
| When adding schedule to a snapshot policy, the count for that schedule
must be specified.

| 1638408
| When adding schedule to a snapshot policy, the schedule name must be
specified.

| 1638413
| Schedule not found.

| 1638417
| Specified policy name is invalid.

| 1638451
| This operation would result in total snapshot count for the policy to
exceed maximum supported count.
```

```
| 1638506
| Schedule already has a prefix. Cannot associate another prefix to the
same schedule.

| 1638508
| Another schedule has the same prefix within this policy.

| 1638526
| This operation is not supported on a node Vserver.

| 1638527
| Policy name already exists.

| 1638528
| This operation is not supported in a mixed-version cluster.

| 1638531
| This operation is not supported because specified policy is owned by the
cluster admin.

| 2621706
| The specified SVM UUID and SVM name refer to different SVMs.

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#copies]
[.api-collapsible-fifth-title]
copies

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.
```

```
|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule. The retention
period value represents a duration and must be specified in the ISO-8601
duration format. The retention period can be in years, months, days,
hours, and minutes. A period specified for years, months, and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively, for example "P10Y" represents a duration of
10 years. A duration in hours and minutes is represented by
"PT+++<num>+++H" and "PT+++<num>+++M" respectively. The period string must
contain only a single time element that is, either years, months, days,
hours, or minutes. A duration which combines different periods is not
supported, for example "P1Y10M" is not
supported.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
```

```
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

The snapshot policy object is associated with a read-write volume used to
create and delete snapshots at regular intervals.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|A comment associated with the snapshot policy.


|copies
|array[link:#copies[copies]]
a|

|enabled
|boolean
a|Is the snapshot policy enabled?


|name
|string
a|Name of the snapshot policy.


|scope
```

```
|string
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID5fac86579cd14b09720cf6594b44ffe8]]
= Delete a snapshot policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snapshot-policies/{uuid}`#

*Introduced In:* 9.6

Deletes a snapshot policy

== Related ONTAP commands

* `snapshot policy delete`

== Learn more
```

```
* xref:{relative_path}storage_snapshot-policies_endpoint_overview.html[DOC
/storage/snapshot-policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Snapshot policy UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Code

|===
| Error Code | Description

| 1638415
| Cannot delete policy. Reason: Policy is in use by at least one volume.

| 1638416
| Cannot delete policy. Reason: Cannot verify whether policy is in use.

| 1638430
| Cannot delete policy. Reason: Policy is in use by at least one Vserver.
```

```
| 1638431
| Cannot delete built-in policy.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6f77e8e3f861dcf486dc4558fcccd532]]
= Retrieve snapshot policy details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snapshot-policies/{uuid}`#

*Introduced In:* 9.6

Retrieves details of a specific snapshot policy.

== Related ONTAP commands

* `snapshot policy show`

== Learn more

* xref:{relative_path}storage_snapshot-policies_endpoint_overview.html[DOC
/storage/snapshot-policies]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
```

```
|True
a|Snapshot policy UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A comment associated with the snapshot policy.


|copies
|array[link:#copies[copies]]
a|

|enabled
|boolean
a|Is the snapshot policy enabled?


|name
|string
a|Name of the snapshot policy.


|scope
```

```
|string
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.


|uuid
|string
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "retention_period": "string",
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
```

```
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
```

```
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#schedule]
```

```
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#copies]
[.api-collapsible-fifth-title]
copies

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
```

```
|string
a|The retention period of snapshots for this schedule. The retention
period value represents a duration and must be specified in the ISO-8601
duration format. The retention period can be in years, months, days,
hours, and minutes. A period specified for years, months, and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively, for example "P10Y" represents a duration of
10 years. A duration in hours and minutes is represented by
"PT+++<num>+++H" and "PT+++<num>+++M" respectively. The period string must
contain only a single time element that is, either years, months, days,
hours, or minutes. A duration which combines different periods is not
supported, for example "P1Y10M" is not
supported.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.
```

```
|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID556593208f681281746d850c44f7825b]]
= Update a snapshot policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/snapshot-policies/{uuid}`#

*Introduced In:* 9.6

Updates a snapshot policy

== Related ONTAP commands

* `snapshot policy modify`
* `snapshot policy modify-schedule`
* `snapshot policy add-schedule`

== Learn more

* xref:{relative_path}storage_snapshot-policies_endpoint_overview.html[DOC
/storage/snapshot-policies]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Snapshot policy UUID

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|A comment associated with the snapshot policy.


|enabled
|boolean
a|Is the snapshot policy enabled?


|scope
|string
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|uuid
|string
```

```
a|

|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "enabled": 1,
  "scope": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Code

|===
| Error Code | Description

| 1638414
| Cannot enable policy. Reason: Schedule not found.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#copies]
[.api-collapsible-fifth-title]
copies

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule. The retention
period value represents a duration and must be specified in the ISO-8601
duration format. The retention period can be in years, months, days,
hours, and minutes. A period specified for years, months, and days is
represented in the ISO-8601 format as "P+++<num>+++Y", "P+++<num>+++M",
"P+++<num>+++D" respectively, for example "P10Y" represents a duration of
10 years. A duration in hours and minutes is represented by
"PT+++<num>+++H" and "PT+++<num>+++M" respectively. The period string must
contain only a single time element that is, either years, months, days,
hours, or minutes. A duration which combines different periods is not
supported, for example "P1Y10M" is not
supported.+++</num>++++++</num>++++++</num>++++++</num>++++++</num>+++


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM. This field cannot be specified in a PATCH method.


|uuid
|string
a|The unique identifier of the SVM. This field cannot be specified in a
PATCH method.


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

The snapshot policy object is associated with a read-write volume used to
create and delete snapshots at regular intervals.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|A comment associated with the snapshot policy.


|enabled
|boolean
a|Is the snapshot policy enabled?


|scope
|string
```

```
a|Set to "svm" when the request is on a data SVM, otherwise set to
"cluster".


|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage snapshot policies and schedules

:leveloffset: +1


[[ID46d84c9acc4ab467898eefa483719e42]]
= Manage snapshot policies and schedules



== Overview

In ONTAP, scheduled snapshot creation works based on the schedules
associated with snapshot policies.
ONTAP provides six cluster-wide schedules: "5min", "8hour", "hourly",
"daily", "weekly" and "monthly".
A snapshot policy is created using at least one of these schedules and up
to 5 schedules can be associated with a snapshot policy.
A snapshot policy can be linked to a storage object and based on the
```

```
schedule in the policy, snapshots are created on the object at that
interval.
Each schedule in a snapshot policy has a snapshot name prefix attached to
it. Every snapshot created using this policy has this prefix in its name.
There is also a retention count associated with every schedule. This count
indicates the maximum number of snapshots that can exist for a given
schedule.
Once the snapshot count reaches the retention count, on the next create
operation, the oldest snapshot is deleted.
A retention period can be associated with every schedule. During snapshot
creation, this period is set as SnapLock expiry time on snapshot locking
enabled volumes.

A schedule can be added, modified or deleted from a snapshot policy.

== Snapshot policy schedule APIs

The following APIs are used to perform operations related to snapshot
policy schedules:

&ndash; POST      /api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/

&ndash; GET       /api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/

&ndash; GET       /api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

&ndash; PATCH     /api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

&ndash; DELETE    /api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

== Examples

=== Adding schedule to a snapshot policy

The POST operation is used to create a schedule for a snapshot policy with
the specified attributes.

----

# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/
```

```
# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json"
-d '{"schedule.uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c", "count":
"5", "prefix": "new_hourly" }'

# The response:
HTTP/1.1 201 Created
Date: Wed, 29 May 2019 22:41:33 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-
005056bbab9c/schedules
Content-Length: 271
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "snapshot_policy": {
      "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
    },
    "schedule": {
      "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c"
    },
    "count": 5,
    "prefix": "new_monthly"
  }
]
}
----

=== Retrieving snapshot policy schedules

The GET operation is used to retrieve snapshot policy schedules.

----

# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 22:49:58 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 898
Content-Type: application/json
{
"records": [
  {
    "snapshot_policy": {
      "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
    },
    "schedule": {
      "uuid": "63d017dc-818a-11e9-b4f4-005056bbab9c",
      "name": "5min"
    }
  },
  {
    "snapshot_policy": {
      "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
    },
    "schedule": {
      "uuid": "64a5c5da-818a-11e9-b4f4-005056bbab9c",
      "name": "8hour"
    }
  },
  {
    "snapshot_policy": {
      "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
    },
    "schedule": {
      "uuid": "63e21a3e-818a-11e9-b4f4-005056bbab9c",
      "name": "daily"
    }
  },
  {
    "snapshot_policy": {
      "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
    },
    "schedule": {
      "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",
      "name": "monthly"
    }
  }
```

```
],
"num_records": 4
}
----

=== Retrieving the attributes of a specific snapshot policy schedule

The GET operation is used to retrieve the attributes of a specific
snapshot policy schedule.

----

# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 22:54:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 238
Content-Type: application/json
{
"snapshot_policy": {
  "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
},
"schedule": {
  "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",
  "name": "monthly"
},
"count": 5,
"prefix": "new_monthly",
"snapmirror_label": "-",
"retention_period": "PT20M"
}
----

=== Updating a snapshot policy schedule
```

```
The PATCH operation is used to update the specific attributes of a
snapshot policy.

----

# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -d '{"count": "10" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:08:00 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
----

=== Deleting a snapshot policy

The DELETE operation is used to delete a snapshot policy.

----

# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:12:32 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

```
----




[[IDf7c154e432575535af7f9a33ed2893ab]]
= Retrieve snapshot policy schedules

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snapshot-policies/{snapshot_policy.uuid}/schedules`#

*Introduced In:* 9.8

Retrieves a collection of snapshot policy schedules.

== Related ONTAP commands

* `snapshot policy show`

== Learn more

* xref:{relative_path}storage_snapshot-
policies_snapshot_policy.uuid_schedules_endpoint_overview.html[DOC
/storage/snapshot-policies/{snapshot_policy.uuid}/schedules]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|snapshot_policy.uuid
|string
|path
|True
a|Snapshot policy UUID


|schedule.name
|string
```

```
|query
|False
a|Filter by schedule.name


|schedule.uuid
|string
|query
|False
a|Filter by schedule.uuid


|retention_period
|string
|query
|False
a|Filter by retention_period

* Introduced in: 9.12


|snapshot_policy.name
|string
|query
|False
a|Filter by snapshot_policy.name


|prefix
|string
|query
|False
a|Filter by prefix


|count
|integer
|query
|False
a|Filter by count


|snapmirror_label
|string
|query
|False
a|Filter by snapmirror_label
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#snapshot_policy_schedule[snapshot_policy_schedule]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "prefix": "string",
      "retention_period": "string",
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snapmirror_label": "string",
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the snapshot policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#snapshot_policy_schedule]
[.api-collapsible-fifth-title]
snapshot_policy_schedule

The snapshot policy schedule object is associated with a snapshot policy
```

```
and it defines the interval at which snapshots are created and deleted.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
```

```
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
```

```
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID7e26ffdfa4a782a56c8419e5c63b69a1]]
= Add a schedule to a snapshot policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/snapshot-policies/{snapshot_policy.uuid}/schedules`#

*Introduced In:* 9.8

Adds a schedule to a snapshot policy.

== Required properties

* `schedule.uuid` or `schedule.name` - Schedule at which snapshots are
captured on the volume.
* `count` - Number of snapshots to maintain for this schedule.

== Recommended optional properties

* `prefix` - Prefix to use when creating snapshots at regular intervals.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `prefix` - Value of `schedule.name`

== Related ONTAP commands

* `snapshot policy add-schedule`

== Learn more

* xref:{relative_path}storage_snapshot-
policies_snapshot_policy.uuid_schedules_endpoint_overview.html[DOC
```

```
/storage/snapshot-policies/{snapshot_policy.uuid}/schedules]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|snapshot_policy.uuid
|string
|path
|True
a|Snapshot policy UUID


|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
```

```
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "prefix": "string",
  "retention_period": "string",
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 201, Created

```
=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1638407
| When adding schedule to a snapshot policy, the count for that schedule
must be specified.

| 1638410
| Specified schedule already exists in snapshot policy.

| 1638413
| Schedule not found.

| 1638451
| This operation would result in total snapshot count for the policy to
exceed maximum supported count.

| 1638508
| Another schedule has the same prefix within this policy.
```

```
| 1638510
| Duplicate prefix.

| 1638528
| This operation is not supported in a mixed-version cluster.

| 1638531
| This operation is not supported because specified policy is owned by the
cluster admin.

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the snapshot policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#snapshot_policy_schedule]
[.api-collapsible-fifth-title]
snapshot_policy_schedule
```

```
The snapshot policy schedule object is associated with a snapshot policy
and it defines the interval at which snapshots are created and deleted.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.
```

```
|===


//end collapsible .Definitions block
====



[[ID7a19275b02a8186fc198d3fafb246098]]
= Delete a schedule from a snapshot policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}`#

*Introduced In:* 9.8

Deletes a schedule from a snapshot policy

== Related ONTAP commands

* `snapshot policy remove-schedule`

== Learn more

* xref:{relative_path}storage_snapshot-
policies_snapshot_policy.uuid_schedules_endpoint_overview.html[DOC
/storage/snapshot-policies/{snapshot_policy.uuid}/schedules]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|snapshot_policy.uuid
|string
|path
|True
```

```
a|Snapshot policy UUID


|schedule.uuid
|string
|path
|True
a|Snapshot policy schedule UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Code

|===
| Error Code | Description

| 1638411
| At least minimum number of schedule must be specified for the policy.

| 1638412
| Schedule does not exist in snapshot policy.
|===



[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID6d8c7e7084b702329b9bfe112cbc578d]]
= Retrieve snapshot policy schedule details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}`#
```

```
*Introduced In:* 9.8

Retrieves details of a specific snapshot policy schedule.

== Related ONTAP commands

* `snapshot policy show`

== Learn more

* xref:{relative_path}storage_snapshot-
policies_snapshot_policy.uuid_schedules_endpoint_overview.html[DOC
/storage/snapshot-policies/{snapshot_policy.uuid}/schedules]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|snapshot_policy.uuid
|string
|path
|True
a|Snapshot policy UUID


|schedule.uuid
|string
|path
|True
a|Snapshot policy schedule ID


|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|count
|integer
a|The number of snapshots to maintain for this schedule.


|prefix
|string
a|The prefix to use while creating snapshots at regular intervals.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.


|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "prefix": "string",
  "retention_period": "string",
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

```
|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===
```

```
[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the snapshot policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID2d4a15aa8525e99980718975e5ad7a65]]
= Update a snapshot policy schedule

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/snapshot-
```

```
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}`#

*Introduced In:* 9.8

Updates a snapshot policy schedule

== Related ONTAP commands

* `snapshot policy modify-schedule`

== Learn more

* xref:{relative_path}storage_snapshot-
policies_snapshot_policy.uuid_schedules_endpoint_overview.html[DOC
/storage/snapshot-policies/{snapshot_policy.uuid}/schedules]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|snapshot_policy.uuid
|string
|path
|True
a|Snapshot policy UUID


|schedule.uuid
|string
|path
|True
a|Snapshot policy schedule UUID

|===

== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.


|===


.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "retention_period": "string",
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Code

|===
| Error Code | Description

| 1638451
| This operation would result in total snapshot count for the policy to
exceed maximum supported count.

| 918253
| Incorrect format for the retention period, duration must be in the ISO-
8601 format.
|===


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```

```
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the snapshot policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
```

```
a|

|uuid
|string
a|

|===


[#snapshot_policy_schedule]
[.api-collapsible-fifth-title]
snapshot_policy_schedule

The snapshot policy schedule object is associated with a snapshot policy
and it defines the interval at which snapshots are created and deleted.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of snapshots to maintain for this schedule.


|retention_period
|string
a|The retention period of snapshots for this schedule.


|schedule
|link:#schedule[schedule]
a|

|snapmirror_label
|string
a|Label for SnapMirror operations


|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the snapshot policy.
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
```

```
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage switches

:leveloffset: +1


[[ID033f09e5cdc6988086c7d8477ef119c5]]
= Manage storage switches



== Retrieving storage switch information

The storage switch GET API retrieves all of the switches in the cluster.

'''

== Examples

=== 1) Retrieves a list of storage switches from the cluster

==== The following example shows the response with a list of storage
switches in the cluster:

'''

----
```

```
# The API:
/api/storage/switches

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/switches" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "name": "Brocade_10.226.57.206",
    "_links": {
      "self": {
        "href": "/api/storage/switches/Brocade_10.226.57.206"
      }
    }
  },
  {
    "name": "Brocade_10.226.57.207",
    "_links": {
      "self": {
        "href": "/api/storage/switches/Brocade_10.226.57.207"
      }
    }
  },
  {
    "name": "Brocade_10.226.57.208",
    "_links": {
      "self": {
        "href": "/api/storage/switches/Brocade_10.226.57.208"
      }
    }
  },
  {
    "name": "Brocade_10.226.57.209",
    "_links": {
      "self": {
        "href": "/api/storage/switches/Brocade_10.226.57.209"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
```

```
    "href": "/api/storage/switches/"
  }
}
}
----

'''

=== 2) Retrieves a specific storage switch from the cluster

==== The following example shows the response of the requested storage
switch. If there is no storage switch with the requested name, an error is
returned.

'''

----

# The API:
/api/storage/switches/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/switches/Brocade_10.226.57.206"
-H "accept: application/hal+json"

# The response:
{
"name": "Brocade_10.226.57.206",
"domain_id": 5,
"switch_fabric_name": "100050eb1a238892",
"fw_version": "v7.2.1c1",
"ip_address": "10.226.57.206",
"is_director_class": false,
"local": false,
"monitoring_enabled": true,
"model": "Brocade6510",
"role": "subordinate",
"state": "ok",
"symbolic_name": "rtp-fc01-41kk11",
"vendor": "brocade",
"wwn": "100050eb1a1ef7d7",
"power_supply_units": [
  {
    "name": "Power Supply #1",
    "state": "ok"
  },
```

```
    {
      "name": "Power Supply #2",
      "state": "ok"
    }
  ],
  "temperature_sensors": [
    {
      "name": "SLOT #0: TEMP #1",
      "reading": 52,
      "state": "ok"
    }
  ],
  "ports": [
    {
      "name": "FC port 0/0",
      "mode": "f_port",
      "wwn": "200050eb1a1ef7d7",
      "enabled": true,
      "state": "online",
      "speed": 16,
      "sfp": {
        "type": "small_form_factor",
        "transmitter_type": "short_wave_laser",
        "serial_number": "HAA2140310058E5"
      }
    },
    {
      "name": "FC port 0/1",
      "mode": "f_port",
      "wwn": "200050eb1a1ef2d7",
      "enabled": true,
      "state": "online",
      "speed": 16,
      "sfp": {
        "type": "small_form_factor",
        "transmitter_type": "short_wave_laser",
        "serial_number": "HAA2140310058E5"
      }
    },
    {
      "name": "FC port 0/2",
      "mode": "f_port",
      "wwn": "200050eb1a1ef7d0",
      "enabled": true,
      "state": "online",
      "speed": 16,
```

```
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  },
  {
    "name": "FC port 0/3",
    "mode": "f_port",
    "wwn": "200050eb1a1ef7d7",
    "enabled": true,
    "state": "online",
    "speed": 16,
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  },
  {
    "name": "FC port 0/4",
    "mode": "f_port",
    "wwn": "200050eb1a1ef2d7",
    "enabled": true,
    "state": "online",
    "speed": 16,
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  },
  {
    "name": "FC port 0/5",
    "mode": "f_port",
    "wwn": "200050eb1a1ef7d0",
    "enabled": true,
    "state": "online",
    "speed": 16,
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  }
],
```

```
"connections": [
  {
    "source_port": {
      "name": "FC port 0/0",
      "wwn": "200050eb1a236efd",
      "mode": "f_port"
    },
    "peer_port": {
      "wwn": "2100000e1e30ac5f",
      "connection": "sti8020mcc-htp-006:fcvi_device_1",
      "type": "fcvi_adapter",
      "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
    }
  },
  {
    "source_port": {
      "name": "FC port 0/1",
      "wwn": "200150eb1a236efd",
      "mode": "f_port"
    },
    "peer_port": {
      "wwn": "21000024ff72c0c9",
      "connection": "sti8020mcc-htp-006:2b",
      "type": "fcp_adapter",
      "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
    }
  },
  {
    "source_port": {
      "name": "FC port 0/2",
      "wwn": "200250eb1a236efd",
      "mode": "f_port"
    },
    "peer_port": {
      "wwn": "21000024ff72c0cb",
      "connection": "sti8020mcc-htp-006:2d",
      "type": "fcp_adapter",
      "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
    }
  }
],
"fans": [
  {
    "name": "FAN #1",
    "speed": 7336,
    "state": "ok"
```

```
  },
  {
    "name": "FAN #2",
    "speed": 7336,
    "state": "ok"
  }
],
"paths": [
     {
       "adapter": {
         "name": "2a",
         "wwn": "21000024ff6c4bc0",
         "type": "fcp_initiator"
       },
       "port": {
         "name": "FC port 0/4",
         "speed": 8
       },
       "node": {
         "name": "sti8020mcc-htp-005",
         "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
         "_links": {
           "self": {
             "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-
00a0985bd455"
           }
         }
       },
       "_links": {
         "self": {
           "href": "/api/storage/ports/382cb083-4416-11eb-ad1d-
00a0985bd455/2a"
         }
       }
     },
     {
       "adapter": {
         "name": "2c",
         "wwn": "21000024ff6c4bc2",
         "type": "fcp_initiator"
       },
       "port": {
         "name": "FC port 0/5",
         "speed": 8
       },
       "node": {
```

```
          "name": "sti8020mcc-htp-005",
          "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-
00a0985bd455"
            }
          }
        },
        "_links": {
          "self": {
            "href": "/api/storage/ports/382cb083-4416-11eb-ad1d-
00a0985bd455/2c"
          }
        }
      },
      {
        "adapter": {
          "name": "fcvi_device_0",
          "wwn": "2100000e1e09d5d2",
          "type": "fc_vi"
        },
        "port": {
          "name": "FC port 0/3",
          "speed": 16
        },
        "node": {
          "name": "sti8020mcc-htp-005",
          "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-
00a0985bd455"
            }
          }
        }
      },
      {
        "adapter": {
          "name": "2a",
          "wwn": "21000024ff72c0c8",
          "type": "fcp_initiator"
        },
        "port": {
          "name": "FC port 0/1",
          "speed": 8
```

```
        },
        "node": {
          "name": "sti8020mcc-htp-006",
          "uuid": "364fbba8-4416-11eb-8e72-00a098431045",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/364fbba8-4416-11eb-8e72-
00a098431045"
            }
          }
        },
        "_links": {
          "self": {
            "href": "/api/storage/ports/364fbba8-4416-11eb-8e72-
00a098431045/2a"
          }
        }
      },
      {
        "adapter": {
          "name": "2c",
          "wwn": "21000024ff72c0ca",
          "type": "fcp_initiator"
        },
        "port": {
        "name": "FC port 0/2",
        "speed": 8
        },
        "node": {
          "name": "sti8020mcc-htp-006",
          "uuid": "364fbba8-4416-11eb-8e72-00a098431045",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/364fbba8-4416-11eb-8e72-
00a098431045"
            }
          }
        },
        "_links": {
          "self": {
            "href": "/api/storage/ports/364fbba8-4416-11eb-8e72-
00a098431045/2c"
          }
        }
      }
],
```

```
"_links": {
  "self": {
    "href": "/api/storage/switches/Brocade_10.226.57.206"
  }
}
}
----

'''




[[ID71e2e41a9a64da27aaaa2849b049c5b5]]
= Retrieve storage switches

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/switches`#

*Introduced In:* 9.9

Retrieves a collection of storage switches.

== Related ONTAP commands

* `storage switch show`

== Learn more

* xref:{relative_path}storage_switches_endpoint_overview.html[DOC
/storage/switches]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|fields
|array[string]
```

```
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#storage_switch[storage_switch]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "connections": [
        {
          "peer_port": {
            "connection": "string",
            "type": "string",
            "unique_id": "string",
```

```
            "wwn": "string"
          },
          "source_port": {
            "mode": "string",
            "name": "string",
            "wwn": "string"
          }
        }
      ],
      "errors": [
        {
          "component": {
            "name": "string"
          },
          "reason": {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          },
          "severity": "string",
          "type": "string"
        }
      ],
      "fabric_name": "string",
      "fans": [
        {
          "name": "string",
          "state": "string"
        }
      ],
      "firmware_version": "string",
      "ip_address": "string",
      "model": "string",
      "monitored_blades": [
        "integer"
      ],
      "name": "string",
      "paths": [
        {
          "adapter": {
            "name": "string",
```

```
            "type": "string",
            "wwn": "string"
          },
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "port": {
            "name": "string"
          }
        }
      ],
      "ports": [
        {
          "mode": "string",
          "name": "string",
          "sfp": {
            "serial_number": "string",
            "transmitter_type": "string",
            "type": "string"
          },
          "state": "string",
          "wwn": "string"
        }
      ],
      "power_supply_units": [
        {
          "name": "string",
          "state": "string"
        }
      ],
      "role": "string",
      "state": "string",
      "symbolic_name": "string",
      "temperature_sensors": [
        {
          "name": "string",
          "state": "string"
        }
      ],
      "vendor": "string",
```

```
      "vsans": [
        {
          "load_balancing_types": "string",
          "name": "string",
          "state": "string"
        }
      ],
      "wwn": "string",
      "zones": [
        {
          "name": "string",
          "port": {
            "id": "string",
            "name": "string"
          },
          "wwn": "string"
        }
      ]
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#peer_port]
[.api-collapsible-fifth-title]
peer_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|connection
|string
a|Storage switch peer port host and name


|type
|string
a|Storage switch peer type


|unique_id
|string
a|Storage switch peer unique ID


|wwn
|string
a|Storage switch peer port world wide name


|===


[#source_port]
```

```
[.api-collapsible-fifth-title]
source_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mode
|string
a|Storage switch port operating mode


|name
|string
a|Storage switch port name


|wwn
|string
a|Storage switch peer port world wide name


|===


[#connections]
[.api-collapsible-fifth-title]
connections

[cols=3*,options=header]
|===
|Name
|Type
|Description

|peer_port
|link:#peer_port[peer_port]
a|

|source_port
|link:#source_port[source_port]
a|

|===
```

```
[#component]
[.api-collapsible-fifth-title]
component

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Error component ID


|name
|string
a|Error component name


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===
```

```
[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|component
|link:#component[component]
a|

|reason
|link:#error[error]
a|

|severity
```

```
|string
a|Error component severity


|type
|string
a|Error component type


|===


[#fans]
[.api-collapsible-fifth-title]
fans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Storage switch fan name


|speed
|integer
a|Storage switch fan speed


|state
|string
a|Storage switch fan state


|===


[#adapter]
[.api-collapsible-fifth-title]
adapter

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|name
|string
a|Node adapter name


|type
|string
a|Node adapter type


|wwn
|string
a|Node adapter world wide name


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===


[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Storage switch port name


|speed
|integer
a|Storage switch port speed, in Gbps


|===


[#paths]
[.api-collapsible-fifth-title]
paths

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|adapter
|link:#adapter[adapter]
a|

|node
|link:#node[node]
a|

|port
|link:#port[port]
a|

|===


[#sfp]
[.api-collapsible-fifth-title]
sfp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|serial_number
|string
a|Storage switch port SFP serial number


|transmitter_type
|string
a|Storage switch port SFP transmitter type


|type
|string
a|Storage switch port SFP type


|===


[#ports]
[.api-collapsible-fifth-title]
```

```
ports

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether the storage switch port is enabled.


|mode
|string
a|Storage switch port mode


|name
|string
a|Storage switch port name


|sfp
|link:#sfp[sfp]
a|

|speed
|integer
a|Storage switch port speed, in Gbps


|state
|string
a|Storage switch port state


|wwn
|string
a|Storage switch port world wide name


|===


[#power_supply_units]
[.api-collapsible-fifth-title]
```

```
power_supply_units

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Power supply unit name


|state
|string
a|Power supply unit state


|===


[#temperature_sensors]
[.api-collapsible-fifth-title]
temperature_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Temperature sensor name


|reading
|integer
a|Temperature sensor reading, in degrees celsius.


|state
|string
a|Temperature sensor state


|===
```

```
[#vsans]
[.api-collapsible-fifth-title]
vsans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Storage switch VSAN ID


|iod
|boolean
a|Indicates whether in-order delivery is set for a zone.


|load_balancing_types
|string
a|Storage switch VSAN load balancing type


|name
|string
a|Storage switch VSAN name


|state
|string
a|Storage switch VSAN Port state


|===


[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|id
|string
a|Storage switch zone port ID


|name
|string
a|Storage switch zone port


|===


[#zones]
[.api-collapsible-fifth-title]
zones

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Storage switch zone ID


|name
|string
a|Storage switch zone name


|port
|link:#port[port]
a|

|wwn
|string
a|Storage switch zone world wide name


|===
```

```
[#storage_switch]
[.api-collapsible-fifth-title]
storage_switch

The Storage switch object describes the storage switch properties,
features and cabling.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|connections
|array[link:#connections[connections]]
a|

|director_class
|boolean
a|

|domain_id
|integer
a|Domain ID


|errors
|array[link:#errors[errors]]
a|

|fabric_name
|string
a|Storage switch fabric name


|fans
|array[link:#fans[fans]]
a|

|firmware_version
|string
a|Storage switch firmware version


|ip_address
```

```
|string
a|IP Address


|local
|boolean
a|Indicates whether the storage switch is directly connected to the
reporting cluster.


|model
|string
a|Storage switch model.


|monitored_blades
|array[integer]
a|Indicates the blades that are being monitored for a director-class
switch.


|monitoring_enabled
|boolean
a|Indicates whether monitoring is enabled for the storage switch.


|name
|string
a|Storage switch name


|paths
|array[link:#paths[paths]]
a|

|ports
|array[link:#ports[ports]]
a|

|power_supply_units
|array[link:#power_supply_units[power_supply_units]]
a|

|role
|string
a|Storage switch role in fabric.
```

```
|state
|string
a|Storage switch state


|symbolic_name
|string
a|Storage switch symbolic name


|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|

|vendor
|string
a|Storage switch vendor


|vsans
|array[link:#vsans[vsans]]
a|

|wwn
|string
a|Storage switch world wide name


|zones
|array[link:#zones[zones]]
a|

|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
```

```
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[ID22266f4005e0f4afbdc27237061e039a]]
= Retrieve a specific storage switch

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/switches/{name}`#

*Introduced In:* 9.9

Retrieves a specific storage switch.

== Related ONTAP commands

* `storage switch show`

== Learn more

* xref:{relative_path}storage_switches_endpoint_overview.html[DOC
/storage/switches]
```

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|connections
|array[link:#connections[connections]]
a|

|director_class
|boolean
a|

|domain_id
|integer
a|Domain ID
```

```
|errors
|array[link:#errors[errors]]
a|

|fabric_name
|string
a|Storage switch fabric name


|fans
|array[link:#fans[fans]]
a|

|firmware_version
|string
a|Storage switch firmware version


|ip_address
|string
a|IP Address


|local
|boolean
a|Indicates whether the storage switch is directly connected to the
reporting cluster.


|model
|string
a|Storage switch model.


|monitored_blades
|array[integer]
a|Indicates the blades that are being monitored for a director-class
switch.


|monitoring_enabled
|boolean
a|Indicates whether monitoring is enabled for the storage switch.
```

```
|name
|string
a|Storage switch name


|paths
|array[link:#paths[paths]]
a|

|ports
|array[link:#ports[ports]]
a|

|power_supply_units
|array[link:#power_supply_units[power_supply_units]]
a|

|role
|string
a|Storage switch role in fabric.


|state
|string
a|Storage switch state


|symbolic_name
|string
a|Storage switch symbolic name


|temperature_sensors
|array[link:#temperature_sensors[temperature_sensors]]
a|

|vendor
|string
a|Storage switch vendor


|vsans
|array[link:#vsans[vsans]]
a|

|wwn
|string
```

```
a|Storage switch world wide name


|zones
|array[link:#zones[zones]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "connections": [
    {
      "peer_port": {
        "connection": "string",
        "type": "string",
        "unique_id": "string",
        "wwn": "string"
      },
      "source_port": {
        "mode": "string",
        "name": "string",
        "wwn": "string"
      }
    }
  ],
  "errors": [
    {
      "component": {
        "name": "string"
      },
      "reason": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
      "severity": "string",
```

```
      "type": "string"
    }
  ],
  "fabric_name": "string",
  "fans": [
    {
      "name": "string",
      "state": "string"
    }
  ],
  "firmware_version": "string",
  "ip_address": "string",
  "model": "string",
  "monitored_blades": [
    "integer"
  ],
  "name": "string",
  "paths": [
    {
      "adapter": {
        "name": "string",
        "type": "string",
        "wwn": "string"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": {
        "name": "string"
      }
    }
  ],
  "ports": [
    {
      "mode": "string",
      "name": "string",
      "sfp": {
        "serial_number": "string",
        "transmitter_type": "string",
        "type": "string"
```

```
      },
      "state": "string",
      "wwn": "string"
    }
  ],
  "power_supply_units": [
    {
      "name": "string",
      "state": "string"
    }
  ],
  "role": "string",
  "state": "string",
  "symbolic_name": "string",
  "temperature_sensors": [
    {
      "name": "string",
      "state": "string"
    }
  ],
  "vendor": "string",
  "vsans": [
    {
      "load_balancing_types": "string",
      "name": "string",
      "state": "string"
    }
  ],
  "wwn": "string",
  "zones": [
    {
      "name": "string",
      "port": {
        "id": "string",
        "name": "string"
      },
      "wwn": "string"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#peer_port]
[.api-collapsible-fifth-title]
peer_port

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|connection
|string
a|Storage switch peer port host and name


|type
|string
a|Storage switch peer type


|unique_id
|string
a|Storage switch peer unique ID


|wwn
|string
a|Storage switch peer port world wide name


|===


[#source_port]
[.api-collapsible-fifth-title]
source_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mode
|string
a|Storage switch port operating mode


|name
|string
a|Storage switch port name


|wwn
```

```
|string
a|Storage switch peer port world wide name


|===


[#connections]
[.api-collapsible-fifth-title]
connections

[cols=3*,options=header]
|===
|Name
|Type
|Description

|peer_port
|link:#peer_port[peer_port]
a|

|source_port
|link:#source_port[source_port]
a|

|===


[#component]
[.api-collapsible-fifth-title]
component

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Error component ID


|name
|string
a|Error component name
```

```
|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code
```

```
|message
|string
a|Error message


|===


[#errors]
[.api-collapsible-fifth-title]
errors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|component
|link:#component[component]
a|

|reason
|link:#error[error]
a|

|severity
|string
a|Error component severity


|type
|string
a|Error component type


|===


[#fans]
[.api-collapsible-fifth-title]
fans

[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|name
|string
a|Storage switch fan name


|speed
|integer
a|Storage switch fan speed


|state
|string
a|Storage switch fan state


|===


[#adapter]
[.api-collapsible-fifth-title]
adapter

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Node adapter name


|type
|string
a|Node adapter type


|wwn
|string
a|Node adapter world wide name


|===
```

```
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
```

```
a|

|name
|string
a|

|uuid
|string
a|

|===


[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Storage switch port name


|speed
|integer
a|Storage switch port speed, in Gbps


|===


[#paths]
[.api-collapsible-fifth-title]
paths

[cols=3*,options=header]
|===
|Name
|Type
|Description

|adapter
```

```
|link:#adapter[adapter]
a|

|node
|link:#node[node]
a|

|port
|link:#port[port]
a|

|===


[#sfp]
[.api-collapsible-fifth-title]
sfp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|serial_number
|string
a|Storage switch port SFP serial number


|transmitter_type
|string
a|Storage switch port SFP transmitter type


|type
|string
a|Storage switch port SFP type


|===


[#ports]
[.api-collapsible-fifth-title]
ports

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|enabled
|boolean
a|Indicates whether the storage switch port is enabled.


|mode
|string
a|Storage switch port mode


|name
|string
a|Storage switch port name


|sfp
|link:#sfp[sfp]
a|

|speed
|integer
a|Storage switch port speed, in Gbps


|state
|string
a|Storage switch port state


|wwn
|string
a|Storage switch port world wide name


|===


[#power_supply_units]
[.api-collapsible-fifth-title]
power_supply_units

[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|name
|string
a|Power supply unit name


|state
|string
a|Power supply unit state


|===


[#temperature_sensors]
[.api-collapsible-fifth-title]
temperature_sensors

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Temperature sensor name


|reading
|integer
a|Temperature sensor reading, in degrees celsius.


|state
|string
a|Temperature sensor state


|===


[#vsans]
```

```
[.api-collapsible-fifth-title]
vsans

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Storage switch VSAN ID


|iod
|boolean
a|Indicates whether in-order delivery is set for a zone.


|load_balancing_types
|string
a|Storage switch VSAN load balancing type


|name
|string
a|Storage switch VSAN name


|state
|string
a|Storage switch VSAN Port state


|===


[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|id
|string
a|Storage switch zone port ID


|name
|string
a|Storage switch zone port


|===


[#zones]
[.api-collapsible-fifth-title]
zones

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Storage switch zone ID


|name
|string
a|Storage switch zone name


|port
|link:#port[port]
a|

|wwn
|string
a|Storage switch zone world wide name


|===


[#returned_error]
[.api-collapsible-fifth-title]
```

```
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage storage tape devices

:leveloffset: +1


[[ID669a41eab44612b0012581aba4d2133c]]
= Manage storage tape devices
```

```
== Retrieving storage tape information

The storage tape GET API retrieves all of the tapes in the cluster.

'''

== Examples

=== 1) Retrieving a list of tapes from the cluster

The following example returns the list of tapes in the cluster:

'''

----

# The API:
/api/storage/tape-devices

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/tape-devices" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "2d.0",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/2d.0"
      }
    }
  },
```

```
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "2d.0L1",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/2d.0L1"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "qeg-tape-brocade2-8g:0.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/qeg-tape-brocade2-8g%3A0.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
```

```
        }
      }
    },
    "device_id": "stsw-broc6510-01:11.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/stsw-broc6510-01%3A11.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:15.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/stsw-broc6510-01%3A15.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:15.126L1",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/stsw-broc6510-01%3A15.126L1"
```

```
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:22.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/stsw-broc6510-01%3A22.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:23.126",
   "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/stsw-broc6510-01%3A23.126"
      }
    }
  },
],
"num_records": 7,
"_links": {
  "self": {
    "href": "/api/storage/tape-devices"
```

```
  }
}
}
----

'''

=== 2) Retrieving a specific tape device from the cluster

The following example returns the requested tape device. If there is no
tape with the requested UID, an error is returned.

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json"

# The response:
{
"node": {
  "uuid": "5f5275eb-5315-11eb-8ac4-00a0985e0dcf",
  "name": "st-8020-1-02",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/5f5275eb-5315-11eb-8ac4-00a0985e0dcf"
    }
  }
},
"device_id": "2d.0",
"type": "tape",
"description": "IBM LTO-6 ULT3580",
"alias": {
  "name": "st7",
  "mapping": "SN[1068000245]"
},
"aliases": [
  {
    "name": "st7",
    "mapping": "SN[1068000245]"
  }
```

```
],
"wwnn": "5001697722ee0010",
"wwpn": "5001697722ee0011",
"serial_number": "1068000245",
"interface": "sas",
"device_state": "offline",
"formats": [
  "LTO-4/5 Native Density",
  "LTO-4/5 Compressed",
  "LTO-6 2.5TB",
  "LTO-6 6.25TB Compressed"
],
"storage_port": {
  "name": "2d"
},
"file_number": -1,
"block_number": -1,
"residual_count": 0,
"density": "low",
"device_names": [
  {
    "rewind_device": "rst0l",
    "no_rewind_device": "nrst0l",
    "unload_reload_device": "urst0l"
  },
  {
    "rewind_device": "rst0m",
    "no_rewind_device": "nrst0m",
    "unload_reload_device": "urst0m"
  },
  {
    "rewind_device": "rst0h",
    "no_rewind_device": "nrst0h",
    "unload_reload_device": "urst0h"
  },
  {
    "rewind_device": "rst0a",
    "no_rewind_device": "nrst0a",
    "unload_reload_device": "urst0a"
  }
],
"reservation_type": "off",
"_links": {
  "self": {
    "href": "/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-
00a0985e0dcf/2d.0"
```

```
  }
}
}
----

'''

== Updating a tape device

The tape PATCH API allows the tape device to be set online or offline,
positioned, and given an alias.

'''

== Examples

=== 1) Taking a tape device offline

The following example takes a tape device offline:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"online": "false"}'

# The response:
{
"job": {
  "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
    }
  }
}
}
----

'''
```

```
=== 2) Bringing a tape device online

The following example brings a tape device online:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"online": "true"}'

# The response:
{
"job": {
  "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
    }
  }
}
}
----

'''

=== 3) Giving a tape device an alias

The following example assigns an alias to a tape device:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"aliases": [{"name": "st0"}]}'
```

```
# The response:
{
}
----

'''

=== 4) Removing a tape device's aliases

The following example clears any aliases previously assigned to a tape
device:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"aliases": []}'

# The response:
{
}
----

'''

=== 5) Rewinding a tape device

The following example rewinds a tape device:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"position": {"operation":
"rewind"}}'
```

```
# The response:
{
"job": {
  "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
    }
  }
}
}
----

'''

=== 6) Forwarding the tape five files

The following example moves the tape forward five file records:

'''

----

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H
"Content-Type: application/hal_json" -d '{"position": {"operation": "fsf",
"count": 5}}'

# The response:
{
"job": {
  "uuid": "954c20d5-306e-11ec-8c2e-00a098b81daa",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/954c20d5-306e-11ec-8c2e-00a098b81daa"
    }
  }
}
}
----
```

```
'''




[[IDce8873b9b783097576fb53ffebf568e3]]
= Retrieve tape devices

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/tape-devices`#

*Introduced In:* 9.9

Retrieves a collection of tape devices.

== Related ONTAP commands

* `storage tape show`

== Learn more

* xref:{relative_path}storage_tape-devices_endpoint_overview.html[DOC
/storage/tape-devices]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|wwnn
|string
|query
|False
a|Filter by wwnn


|alias.mapping
|string
|query
```

```
|False
a|Filter by alias.mapping


|alias.name
|string
|query
|False
a|Filter by alias.name


|residual_count
|integer
|query
|False
a|Filter by residual_count


|aliases.mapping
|string
|query
|False
a|Filter by aliases.mapping

* Introduced in: 9.11


|aliases.name
|string
|query
|False
a|Filter by aliases.name

* Introduced in: 9.11


|reservation_type
|string
|query
|False
a|Filter by reservation_type


|node.name
|string
|query
|False
```

```
a|Filter by node.name


|node.uuid
|string
|query
|False
a|Filter by node.uuid


|type
|string
|query
|False
a|Filter by type


|wwpn
|string
|query
|False
a|Filter by wwpn


|density
|string
|query
|False
a|Filter by density

* Introduced in: 9.11


|description
|string
|query
|False
a|Filter by description


|device_id
|string
|query
|False
a|Filter by device_id
```

```
|device_state
|string
|query
|False
a|Filter by device_state


|interface
|string
|query
|False
a|Filter by interface


|online
|boolean
|query
|False
a|Filter by online

* Introduced in: 9.11


|device_names.rewind_device
|string
|query
|False
a|Filter by device_names.rewind_device


|device_names.unload_reload_device
|string
|query
|False
a|Filter by device_names.unload_reload_device


|device_names.no_rewind_device
|string
|query
|False
a|Filter by device_names.no_rewind_device


|block_number
|integer
|query
```

```
|False
a|Filter by block_number


|formats
|string
|query
|False
a|Filter by formats


|storage_port.name
|string
|query
|False
a|Filter by storage_port.name


|file_number
|integer
|query
|False
a|Filter by file_number


|serial_number
|string
|query
|False
a|Filter by serial_number


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
```

```
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 15
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records


|records
|array[link:#tape_device[tape_device]]
a|

|===


.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "alias": {
        "mapping": "SN[10WT000933]",
        "name": "st6"
      },
      "aliases": [
        {
          "mapping": "SN[10WT000933]",
          "name": "st6"
        }
      ],
      "block_number": 0,
      "density": "low",
      "description": "QUANTUM LTO-8 ULTRIUM",
      "device_id": "1a.0",
      "device_names": [
        {
          "no_rewind_device": "nrst6l",
          "rewind_device": "rst6l",
```

```
          "unload_reload_device": "urst6l"
        }
      ],
      "device_state": "read_write_enabled",
      "file_number": 0,
      "formats": [
        "LTO-7 6TB",
        "LTO-7 15TB Compressed",
        "LTO-8 12TB",
        "LTO-8 30TB Compressed"
      ],
      "interface": "sas",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "position": {
        "count": 5,
        "operation": "rewind"
      },
      "reservation_type": "off",
      "residual_count": 0,
      "serial_number": "10WT00093",
      "storage_port": {
        "name": "2b"
      },
      "type": "tape",
      "wwnn": "500507631295741c",
      "wwpn": "500507631295741c"
    }
  ]
}
====

== Error
```


Status: Default, Error

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===


[#alias]
[.api-collapsible-fifth-title]
alias

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
|string
a|This field will no longer be supported in a future release. Use
aliases.mapping instead.


|name
|string
a|This field will no longer be supported in a future release. Use
```

```
aliases.name instead.


|===


[#aliases]
[.api-collapsible-fifth-title]
aliases

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
|string
a|Alias mapping.


|name
|string
a|Alias name.


|===


[#device_names]
[.api-collapsible-fifth-title]
device_names

[cols=3*,options=header]
|===
|Name
|Type
|Description

|no_rewind_device
|string
a|Device name for no rewind.


|rewind_device
|string
a|Device name for rewind.
```

```
|unload_reload_device
|string
a|Device name for unload or reload operations.


|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
```

```
a|

|===


[#position]
[.api-collapsible-fifth-title]
position

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of times to run position operation.


|operation
|string
a|Position operation.


|===


[#storage_port]
[.api-collapsible-fifth-title]
storage_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Initiator port.


|===


[#tape_device]
```

```
[.api-collapsible-fifth-title]
tape_device

[cols=3*,options=header]
|===
|Name
|Type
|Description

|alias
|link:#alias[alias]
a|

|aliases
|array[link:#aliases[aliases]]
a|

|block_number
|integer
a|Block number.


|density
|string
a|Density.


|description
|string
a|

|device_id
|string
a|

|device_names
|array[link:#device_names[device_names]]
a|

|device_state
|string
a|Operational state of the device.


|file_number
|integer
a|File number.
```

```
|formats
|array[string]
a|Tape cartridge format.


|interface
|string
a|Device interface type.


|node
|link:#node[node]
a|

|online
|boolean
a|

|position
|link:#position[position]
a|

|reservation_type
|string
a|

|residual_count
|integer
a|Residual count of the last I/O operation.


|serial_number
|string
a|

|storage_port
|link:#storage_port[storage_port]
a|

|type
|string
a|Device type.


|wwnn
```

```
|string
a|World Wide Node Name.


|wwpn
|string
a|World Wide Port Name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDbc34fd250dae67cca7c69d53c213931c]]
= Retrieve a tape device

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/tape-devices/{node.uuid}/{device_id}`#

*Introduced In:* 9.9

Retrieves a specific tape device.

== Related ONTAP commands

* `storage tape show`

== Learn more

* xref:{relative_path}storage_tape-devices_endpoint_overview.html[DOC
/storage/tape-devices]


== Parameters
```

```
[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID


|device_id
|string
|path
|True
a|Device ID


|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|alias
|link:#alias[alias]
a|
```

```
|aliases
|array[link:#aliases[aliases]]
a|

|block_number
|integer
a|Block number.


|density
|string
a|Density.


|description
|string
a|

|device_id
|string
a|

|device_names
|array[link:#device_names[device_names]]
a|

|device_state
|string
a|Operational state of the device.


|file_number
|integer
a|File number.


|formats
|array[string]
a|Tape cartridge format.


|interface
|string
a|Device interface type.


|node
```

```
|link:#node[node]
a|

|online
|boolean
a|

|position
|link:#position[position]
a|

|reservation_type
|string
a|

|residual_count
|integer
a|Residual count of the last I/O operation.


|serial_number
|string
a|

|storage_port
|link:#storage_port[storage_port]
a|

|type
|string
a|Device type.


|wwnn
|string
a|World Wide Node Name.


|wwpn
|string
a|World Wide Port Name.


|===


.Example response
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "alias": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "aliases": [
    {
      "mapping": "SN[10WT000933]",
      "name": "st6"
    }
  ],
  "block_number": 0,
  "density": "low",
  "description": "QUANTUM LTO-8 ULTRIUM",
  "device_id": "1a.0",
  "device_names": [
    {
      "no_rewind_device": "nrst6l",
      "rewind_device": "rst6l",
      "unload_reload_device": "urst6l"
    }
  ],
  "device_state": "read_write_enabled",
  "file_number": 0,
  "formats": [
    "LTO-7 6TB",
    "LTO-7 15TB Compressed",
    "LTO-8 12TB",
    "LTO-8 30TB Compressed"
  ],
  "interface": "sas",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "position": {
    "count": 5,
    "operation": "rewind"
```

```
  },
  "reservation_type": "off",
  "residual_count": 0,
  "serial_number": "10WT00093",
  "storage_port": {
    "name": "2b"
  },
  "type": "tape",
  "wwnn": "500507631295741c",
  "wwpn": "500507631295741c"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#returned_error[returned_error]
a|

|===


.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
```

```
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#alias]
[.api-collapsible-fifth-title]
alias

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
|string
a|This field will no longer be supported in a future release. Use
aliases.mapping instead.


|name
|string
a|This field will no longer be supported in a future release. Use
aliases.name instead.


|===


[#aliases]
[.api-collapsible-fifth-title]
aliases

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
```

```
|string
a|Alias mapping.


|name
|string
a|Alias name.


|===


[#device_names]
[.api-collapsible-fifth-title]
device_names

[cols=3*,options=header]
|===
|Name
|Type
|Description

|no_rewind_device
|string
a|Device name for no rewind.


|rewind_device
|string
a|Device name for rewind.


|unload_reload_device
|string
a|Device name for unload or reload operations.


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
```

```
a|

|===


[#position]
[.api-collapsible-fifth-title]
position

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of times to run position operation.


|operation
|string
a|Position operation.


|===


[#storage_port]
[.api-collapsible-fifth-title]
storage_port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Initiator port.


|===


[#error_arguments]
```

```
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



[[IDca803659f3615cf5457a5ef0b4982320]]
= Update a tape device

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/tape-devices/{node.uuid}/{device_id}`#

*Introduced In:* 9.11

Updates a specific tape device.

== Related ONTAP commands

* `storage tape alias-set`
* `storage tape alias-clear`
* `storage tape online`
* `storage tape offline`
* `storage tape position`

== Learn more

* xref:{relative_path}storage_tape-devices_endpoint_overview.html[DOC
/storage/tape-devices]


== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description
```

```
|node.uuid
|string
|path
|True
a|Node UUID


|device_id
|string
|path
|True
a|Device ID


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds.  This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job.  If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 0
* Max value: 120
* Min value: 0

|===

== Request Body



[cols=3*,options=header]
|===
|Name
|Type
|Description

|alias
|link:#alias[alias]
a|
```

```
|aliases
|array[link:#aliases[aliases]]
a|

|block_number
|integer
a|Block number.


|density
|string
a|Density.


|description
|string
a|

|device_id
|string
a|

|device_names
|array[link:#device_names[device_names]]
a|

|device_state
|string
a|Operational state of the device.


|file_number
|integer
a|File number.


|formats
|array[string]
a|Tape cartridge format.


|interface
|string
a|Device interface type.


|node
```

```
|link:#node[node]
a|

|online
|boolean
a|

|position
|link:#position[position]
a|

|reservation_type
|string
a|

|residual_count
|integer
a|Residual count of the last I/O operation.


|serial_number
|string
a|

|storage_port
|link:#storage_port[storage_port]
a|

|type
|string
a|Device type.


|wwnn
|string
a|World Wide Node Name.


|wwpn
|string
a|World Wide Port Name.


|===


.Example request
```

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "alias": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "aliases": [
    {
      "mapping": "SN[10WT000933]",
      "name": "st6"
    }
  ],
  "block_number": 0,
  "density": "low",
  "description": "QUANTUM LTO-8 ULTRIUM",
  "device_id": "1a.0",
  "device_names": [
    {
      "no_rewind_device": "nrst6l",
      "rewind_device": "rst6l",
      "unload_reload_device": "urst6l"
    }
  ],
  "device_state": "read_write_enabled",
  "file_number": 0,
  "formats": [
    "LTO-7 6TB",
    "LTO-7 15TB Compressed",
    "LTO-8 12TB",
    "LTO-8 30TB Compressed"
  ],
  "interface": "sas",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "position": {
    "count": 5,
    "operation": "rewind"
  },
  "reservation_type": "off",
  "residual_count": 0,
  "serial_number": "10WT00093",
  "storage_port": {
```

```
    "name": "2b"
  },
  "type": "tape",
  "wwnn": "500507631295741c",
  "wwpn": "500507631295741c"
}
====

== Response
```

Status: 200, Ok

```
== Response
```

Status: 202, Accepted

```
== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 11403264
| Tape operation "+++<operation>+++\" failed on tape device
\"+++<device_id>+++\". This may indicate tape reached end of data, a
hardware error, an illegal request, an invalid name format, or an aborted
command.+++</device_id>++++++</operation>+++

| 11403265
| Tape operation  "+++<operation>+++\" failed because it encountered an
unexpected file mark on tape device
\"+++<device_id>+++\".+++</device_id>++++++</operation>+++

| 11403266
| Internal error. Unknown tape parameter for tape device
"+++<device_id>+++\".+++</device_id>+++

| 11403267
| Tape device "+++<device_id>+++\" is busy.+++</device_id>+++

| 11403268
```

```
| No tape loaded for "+++<operation>+++\" operation on tape drive
\"+++<device_id>+++\".+++</device_id>++++++</operation>+++

| 11403269
| Tape device "+++<device_id>+++\" unknown.+++</device_id>+++

| 11403270
| Tape alias name "+++<alias>+++\" already exists.+++</alias>+++

| 11403271
| The format of the tape alias name "+++<alias>+++\" is invalid. Use
\"st\" or \"mc\" followed by one or more digits.+++</alias>+++

| 11403273
| Internal error. Tape alias name "+++<alias>+++\" not
saved.+++</alias>+++

| 11403274
| Internal error. Tape alias operation "+++<operation>+++\" for alias name
\"+++<alias>+++\" failed.+++</alias>++++++</operation>+++

| 11403275
| Tape alias name "+++<alias>+++\" not found.+++</alias>+++

| 11403277
| Tape device "+++<device_id>+++\" is reserved by another
host.+++</device_id>+++

| 11403278
| Invalid count value specified with operation "+++<operation>+++\". Valid
values are from 0 to \{1}.+++</operation>+++
|===

Also see the table of common errors in the
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Response_
body[Response body] overview section of this documentation.


== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#alias]
[.api-collapsible-fifth-title]
```

```
alias

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
|string
a|This field will no longer be supported in a future release. Use
aliases.mapping instead.


|name
|string
a|This field will no longer be supported in a future release. Use
aliases.name instead.


|===


[#aliases]
[.api-collapsible-fifth-title]
aliases

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mapping
|string
a|Alias mapping.


|name
|string
a|Alias name.


|===


[#device_names]
```

```
[.api-collapsible-fifth-title]
device_names

[cols=3*,options=header]
|===
|Name
|Type
|Description

|no_rewind_device
|string
a|Device name for no rewind.


|rewind_device
|string
a|Device name for rewind.


|unload_reload_device
|string
a|Device name for unload or reload operations.


|===


[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===


[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===


[#position]
[.api-collapsible-fifth-title]
position

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of times to run position operation.


|operation
|string
a|Position operation.


|===


[#storage_port]
[.api-collapsible-fifth-title]
storage_port
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Initiator port.


|===


[#tape_device]
[.api-collapsible-fifth-title]
tape_device

[cols=3*,options=header]
|===
|Name
|Type
|Description

|alias
|link:#alias[alias]
a|

|aliases
|array[link:#aliases[aliases]]
a|

|block_number
|integer
a|Block number.


|density
|string
a|Density.


|description
|string
a|
```

```
|device_id
|string
a|

|device_names
|array[link:#device_names[device_names]]
a|

|device_state
|string
a|Operational state of the device.


|file_number
|integer
a|File number.


|formats
|array[string]
a|Tape cartridge format.


|interface
|string
a|Device interface type.


|node
|link:#node[node]
a|

|online
|boolean
a|

|position
|link:#position[position]
a|

|reservation_type
|string
a|

|residual_count
|integer
a|Residual count of the last I/O operation.
```

```
|serial_number
|string
a|

|storage_port
|link:#storage_port[storage_port]
a|

|type
|string
a|Device type.


|wwnn
|string
a|World Wide Node Name.


|wwpn
|string
a|World Wide Port Name.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument
```

```
|===


[#returned_error]
[.api-collapsible-fifth-title]
returned_error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====



:leveloffset: -1


= Manage volume efficiency policies
```

```
:leveloffset: +1


[[IDec051ff8655f5d2b69545453e8c63a49]]
= Manage volume efficiency policies



:doctype: book

=== Overview

Volume efficiency policies specify information about efficiency policies
that are applied to the volume.

== Volume efficiency policy APIs

The following APIs are used to perform operations related to volume
efficiency policy information:

&ndash; POST      /api/storage/volume-efficiency-policies

&ndash; GET       /api/storage/volume-efficiency-policies

&ndash; GET       /api/storage/volume-efficiency-policies/\{uuid}

&ndash; PATCH     /api/storage/volume-efficiency-policies/\{uuid}

&ndash; DELETE    /api/storage/volume-efficiency-policies/\{uuid}

== Examples

=== Creating a volume efficiency policy

The POST operation is used to create a volume efficiency policy with the
specified attributes.

----

# The API:
/api/storage/volume-efficiency-policies

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volume-efficiency-policies"
-H "accept: application/hal+json" -d '{"name": "new_policy", "type":
"scheduled", "schedule": { "name": "daily" }, "duration": "2",
"qos_policy": "best_effort", "enabled": "true", "comment": "schedule-
```

```
policy", "svm": { "name": "vs1"}}'

# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volume-efficiency-policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
    "svm": {
      "name": "vs1"
    },
    "name": "new_policy",
    "type": "scheduled",
    "schedule":{
      "name": "daily"
    },
    "duration": "2",
    "qos_policy": "best_effort",
    "enabled": "true",
    "comment": "schedule-policy"
  }
]
}
----

=== Retrieving volume efficiency policy attributes

The GET operation is used to retrieve volume efficiency policy attributes.

----

# The API:
/api/storage/volume-efficiency-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volume-efficiency-policies" -H
"accept: application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:17:17 GMT
Server: libzapid-http
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 686
Content-Type: application/json
{
"records": [
  {
    "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c112527-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c1c1656-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default-1weekly",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c1c1656-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
    "name": "none",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c228b82-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volume-efficiency-policies/"
  }
```

```
}
}
----

=== Retrieving the attributes of a specific volume efficiency policy

The GET operation is used to retrieve the attributes of a specific volume
efficiency policy.

= The API:

/api/storage/volume-efficiency-policies/\{uuid}

= The call:

curl -X GET "https://+++<mgmt-ip>+++/api/storage/volume-efficiency-
policies/3c112527-2fe8-11e9-b55e-005056bbf1c8" -H "accept:
application/hal+json"+++</mgmt-ip>+++

= The response:

HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:24:48 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 381
Content-Type: application/json
{
"uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
"name": "new_policy",
"type": "scheduled",
"schedule": {
  "name": "daily"
}
"duration": "2",
"qos_policy": "best_effort",
"enabled": "true",
"comment": "schedule-policy",
"svm": {
      "name": "vs1"
}
"_links": {
  "self": {
    "href": "/api/storage/volume-efficiency-policies/3c112527-2fe8-11e9-
b55e-005056bbf1c8"
```

```
  }
}
}

----

### Updating a volume efficiency policy
The PATCH operation is used to update the specific attributes of a volume
efficiency policy.
----

= The API:

/api/storage/volume-efficiency-policies/\{uuid}

= The call:

curl -X PATCH  "https://+++<mgmt-ip>+++/api/storage/volume-efficiency-
policies/ae9e65c4-4506-11e9-aa44-005056bbc848" -d '{"duration": "3" }' -H
"accept: application/hal+json"+++</mgmt-ip>+++

= The response:

HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:27:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json

----

### Deleting a volume efficiency policy
The DELETE operation is used to delete a volume efficiency policy.
----

= The API:

/api/storage/volume-efficiency-policies/\{uuid}

= The call:

curl -X DELETE "https://+++<mgmt-ip>+++/api/storage/volume-efficiency-
policies/ ae9e65c4-4506-11e9-aa44-005056bbc848" -H "accept:
application/hal+json"+++</mgmt-ip>+++
```

```
= The response:

HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:19:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

**Retrieve volume efficiency policies**

GET `/storage/volume-efficiency-policies`

**Introduced In:** 9.8

Retrieves a collection of volume efficiency policies.

#### Related ONTAP commands

- `volume efficiency policy show`

#### Learn more

- DOC /storage/volume-efficiency-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| enabled | boolean | query | False | Filter by enabled |
| type | string | query | False | Filter by type |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| schedule.name | string | query | False | Filter by schedule.name |
| name | string | query | False | Filter by name |
| uuid | string | query | False | Filter by uuid |
| duration | integer | query | False | Filter by duration |
| qos_policy | string | query | False | Filter by qos_policy |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| start_threshold_percent | integer | query | False | Filter by start_threshold_percent |
| comment | string | query | False | Filter by comment |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[volume_efficiency_policy] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "name": "string",
      "qos_policy": "string",
      "schedule": {
        "name": "string"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

schedule

| Name | Type | Description |
|---|---|---|
| name | string | Schedule at which volume efficiency policies are captured on the SVM. Some common schedules already defined in the system are hourly, daily, weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency policies with custom schedules can be referenced. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume_efficiency_policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment associated with the volume efficiency policy. |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| name | string | Name of the volume efficiency policy. |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a volume efficiency policy

POST `/storage/volume-efficiency-policies`

**Introduced In:** 9.8

Creates a volume efficiency policy.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the volume efficiency policy.
- `name` - Name for the volume efficiency policy.

#### Recommended optional properties

- `type` - Type of volume policy.
- `schedule` - Schedule the volume efficiency defined in minutes, hourly, daily and weekly.
- `duration` - Indicates the allowed duration for a session for policy type "scheduled".
- `start_threshold_percent` - Indicates the start threshold percentage for the policy type "threshold". It is mutually exclusive of the schedule.
- `qos_policy` - QoS policy for the sis operation.
- `comment` - A comment associated with the volume efficiency policy.
- `enabled` - Is the volume efficiency policy enabled?

#### Default property values

If not specified in POST, the following default property values are assigned:

- `type` - scheduled
- `start_threshold_percent` - 20
- `enabled` - true
- `qos_policy` - best_effort
- `schedule` - daily

#### Related ONTAP commands

- `volume efficiency policy create`

#### Learn more

- DOC /storage/volume-efficiency-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the volume efficiency policy. |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| name | string | Name of the volume efficiency policy. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

**Example request**

```
{
  "comment": "string",
  "name": "string",
  "qos_policy": "string",
  "schedule": {
    "name": "string"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
```

#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 918702 | The specified operation on the volume efficiency policies endpoint is not supported on this platform. |
| 6881341 | Specified schedule not found. |
| 6881344 | Failed to queue specified job. |
| 6881345 | This operation is not permitted on a node SVM. |
| 6881349 | Policy name is not valid. |
| 6881351 | Policy name is too long. |
| 6881352 | Policy name is not valid. |
| 6881362 | Threshold percentage cannot be less than 1 percent. |
| 6881433 | For "threshold" type policy, "schedule and duration" is not supported. |
| 6881435 | Only a policy of type "threshold" can set the "start-threshold-percent" attribute. |
| 6881454 | An efficiency policy of type "threshold" requires an effective cluster version of ONTAP 8.3 or later. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

schedule

| Name | Type | Description |
|---|---|---|
| name | string | Schedule at which volume efficiency policies are captured on the SVM. Some common schedules already defined in the system are hourly, daily, weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency policies with custom schedules can be referenced. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume_efficiency_policy

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the volume efficiency policy. |

| Name | Type | Description |
| --- | --- | --- |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| name | string | Name of the volume efficiency policy. |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a volume efficiency policy

DELETE `/storage/volume-efficiency-policies/{uuid}`

**Introduced In:** 9.8

Deletes a volume efficiency policy.

#### Related ONTAP commands

- `volume efficiency policy delete`

#### Learn more

- DOC /storage/volume-efficiency-policies

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | Volume efficiency policy UUID |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
|---|---|
| 918702 | The specified operation on the volume efficiency policies endpoint is not supported on this platform. |
| 6881346 | The policy was not deleted because the policy is in use by at least one volume. |
| 6881431 | The specified policy is a predefined policy and cannot be deleted. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve volume efficiency policy details

GET `/storage/volume-efficiency-policies/{uuid}`

**Introduced In:** 9.8

Retrieves the details of the specified volume efficiency policy.

#### Related ONTAP commands

- `volume efficiency policy show`

#### Learn more

- DOC /storage/volume-efficiency-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Volume efficiency policy UUID |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment associated with the volume efficiency policy. |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| name | string | Name of the volume efficiency policy. |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "name": "string",
  "qos_policy": "string",
  "schedule": {
    "name": "string"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

schedule

| Name | Type | Description |
|---|---|---|
| name | string | Schedule at which volume efficiency policies are captured on the SVM. Some common schedules already defined in the system are hourly, daily, weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency policies with custom schedules can be referenced. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update a volume efficiency policy

PATCH `/storage/volume-efficiency-policies/{uuid}`

**Introduced In:** 9.8

Updates a volume efficiency policy.

#### Related ONTAP commands

- `volume efficiency policy modify`

#### Learn more

- DOC /storage/volume-efficiency-policies

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Volume efficiency policy UUID |

#### Request Body

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the volume efficiency policy. |

| Name | Type | Description |
| --- | --- | --- |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

**Example request**

```
{
  "comment": "string",
  "qos_policy": "string",
  "schedule": {
    "name": "string"
  },
  "type": "string",
  "uuid": "string"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 918702 | The specified operation on the volume efficiency policies endpoint is not supported on this platform. |
| 6881341 | Specified schedule not found. |
| 6881344 | Failed to queue specified job. |
| 6881351 | Policy name is too long. |
| 6881352 | Policy name is not valid. |
| 6881362 | Threshold percentage cannot be less than 1 percent. |
| 6881433 | For "threshold" type policy, "schedule and duration" is not supported. |
| 6881435 | Only a policy of type "threshold" can set the "start-threshold-percent" attribute. |
| 6881438 | For "inline-only and auto" policy, modification of attributes is not allowed. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

schedule

| Name | Type | Description |
|---|---|---|
| name | string | Schedule at which volume efficiency policies are captured on the SVM. Some common schedules already defined in the system are hourly, daily, weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency policies with custom schedules can be referenced. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume_efficiency_policy

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the volume efficiency policy. |

| Name | Type | Description |
| --- | --- | --- |
| duration | integer | This field is used with the policy type "scheduled" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0. |
| enabled | boolean | Is the volume efficiency policy enabled? |
| qos_policy | string | QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations. |
| schedule | schedule | |
| start_threshold_percent | integer | This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule. |
| type | string | Type of volume efficiency policy. |
| uuid | string | Unique identifier of volume efficiency policy. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage storage volumes

### Manage storage volumes

#### Overview

FlexVol volumes are logical containers used by ONTAP to serve data to clients. They contain file systems in a NAS environment and LUNs in a SAN environment.

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup volume contains several constituents that automatically and transparently share the traffic.

FlexClone volumes are writable, point-in-time copies of a FlexVol volume. At this time, FlexClones of FlexGroups are not supported.

Volumes with SnapLock type Compliance or Enterprise, are referred to as SnapLock volumes. It is possible to create a SnapLock volume by specifying SnapLock parameters.

ONTAP storage APIs allow you to create, modify, and monitor volumes and aggregates.

#### Storage efficiency

Storage efficiency is used to remove duplicate blocks in the data and to compress the data. Efficiency has deduplication, compression, cross volume deduplication, compaction, policy-name, enabled, application_io_size, compression_type and storage_efficiency_mode options. On All Flash systems, all efficiencies are enabled by default, on volume creation. Options such as "background/inline/both" are treated as both, which means both background and inline are enabled for any efficiency option. The option "none" disables both background and inline efficiency. Application-io-size and compression-type decides type of compression behavior in the system. Storage efficiency mode decides if the system is to run in default/efficient mode. Detailed information about each field is available under efficiency object for storage efficiency fields.

To enable any efficiency option on all-flash or FAS systems, background deduplication is always enabled.

#### Quotas

Quotas provide a way to restrict or track the files and space usage by a user, group, or qtree. Quotas are enabled for a specific FlexVol or a FlexGroup volume.

The following APIs can be used to enable or disable and obtain quota state for a FlexVol or a FlexGroup volume:

– PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"true"}'

– PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"false"}'

– GET /api/storage/volumes/{uuid}/?fields=quota.state

#### File System Analytics

File system analytics (FSA) provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. For more information on FSA, see DOC /storage/volumes{volume.uuid}/files/{path} . Analytics can be enabled or disabled on individual volumes.

The following APIs can be used to enable or disable and obtain analytics state for a FlexVol volume or a FlexGroup volume:

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"on"}'

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"off"}'

– GET /api/storage/volumes/{uuid}/?fields=analytics

If the `analytics.state` field is "initializing" or "initialization_paused" and is set to "off", the FSA initialization scan is cancelled. If FSA is turned on again, the initialization scan restarts.

#### QoS

QoS policy and settings enforce Service Level Objectives (SLO) on a volume. SLO can be set by specifying qos.max_throughput_iops and/or qos.max_throughput_mbps or qos.min_throughput_iops and/or qos.min_throughput_mbps. Specifying min_throughput_iops or min_throughput_mbps is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying qos.name or qos.uuid property.

#### Performance monitoring

Performance of a volume can be monitored by the `metric.*` and `statistics.*` fields. These show the performance of the volume in terms of IOPS, latency and throughput. The `metric.*` fields denote an average whereas `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

#### Rebalancing

Non-disruptive capacity rebalancing of a FlexGroup volume is configured by the `rebalancing.*` fields. If not explicitly set, default values are provided. To initiate a capacity rebalancing operation, `rebalancing.state` is set to 'starting'. The `rebalancing.max_runtime` can be optionally set, which is the maximum length of time you want the capacity rebalancing to run for. You can stop capacity rebalancing by setting `rebalancing.state` to 'stopping'. You can also modify the configurations `rebalancing.max_runtime`, `rebalancing.max_threshold`, `rebalancing.min_threshold`, `rebalancing.max_file_moves`, `rebalancing.min_file_size`, and `rebalancing.exclude_snapshots`. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current configuration values. Once the operation is started, changes to the configuration are not allowed, until the capacity rebalanding operations stops, either by exceeding their maximum runtime or by being stopped. To see runtime information about each constituent, for a running rebalancing operation, use the 'rebalancing.engine.*' fields.

#### Volume APIs

The following APIs are used to perform operations related with FlexVol volumes and FlexGroup volumes:

– POST /api/storage/volumes

– GET /api/storage/volumes

– GET /api/storage/volumes/{uuid}

– PATCH /api/storage/volumes/{uuid}

– DELETE /api/storage/volumes/{uuid}

#### Examples

### Creating a volume

The POST request is used to create a new volume and to specify its properties.

```
# The API:
/api/storage/volumes

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name": "vol1",
"aggregates":[{"name":"aggr1"}], "svm":{"name" : "vs1"}}'

# The response:
{
"job": {
  "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
    }
  }
}
}
```

### Creating a SnapLock volume and specifying its properties using POST

```
# The API:
/api/storage/volumes

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name": "vol1",  "aggregates":[{"name":
"aggr1"}],  "svm":{"name" : "vs1"}, "snaplock":{"retention":{"default":
"P20Y"}, "type": "compliance"}}'

# The response:
{
"job": {
  "uuid": "e45b123b-c228-11e8-aa20-0050568e36bb",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e45b123b-c228-11e8-aa20-0050568e36bb"
    }
  }
}
}
```

**Creating a FlexGroup volume and specifying its properties using POST**

```
# The API:
/api/storage/volumes

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name" : "vol1", "state" : "online", "type" :
"RW", "aggregates" : [{"name" : "aggr1"}, {"name" : "aggr2"},
{"name":"aggr3"}], "constituents_per_aggregate" : "1", "svm" : {"name" :
"vs1"}, "size" : "240MB", "encryption" : {"enabled" : "False"},
"efficiency" : {"compression" : "both"}, "autosize" : {"maximum" :
"500MB", "minimum" : "240MB"}}'

# The response:
{
"job": {
  "uuid": "3cfa38bd-3a78-11e9-ae39-0050568ed7dd",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3cfa38bd-3a78-11e9-ae39-0050568ed7dd"
    }
  }
}
}
```

**Creating a FlexGroup volume and specifying its properties using POST when the Performance_NAS license is installed.**

```
# The API:
/api/storage/volumes

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name" : "vol1", "state" : "online", "type" :
"RW", "svm" : {"name" : "vs1"}, "size" : "240TB", "encryption" :
{"enabled" : "False"}, "efficiency" : {"compression" : "both"}, "autosize"
: {"maximum" : "500TB", "minimum" : "240TB"}}'

# The response:
{
"job": {
  "uuid": "3cfa38bd-3a78-11e9-ae39-0050568ed7dd",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3cfa38bd-3a78-11e9-ae39-0050568ed7dd"
    }
  }
}
}
```

**Creating a FlexClone and specifying its properties using POST**

```
# The API:
/api/storage/volumes

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name":"vol1_clone", "clone": {"parent_volume":
{"name": "vol1"}, "is_flexclone": "true"},"svm":{"name": "vs0"}}'

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 26 Feb 2019 09:06:22 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/?name=vol1_clone
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "c9ee0040-39a5-11e9-9b24-00a098439a83",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c9ee0040-39a5-11e9-9b24-00a098439a83"
    }
  }
}
}
```

#### Volumes reported in the GET REST API

**The following types of volumes are reported:**

– RW, DP and LS volumes

– FlexGroup volume

– FlexCache volume

– FlexClone volume

– FlexGroup constituent

**The following volumes are not reported:**

– DEL and TMP type volume

– Node Root volume

– System Vserver volume

– FlexCache constituent

#### Examples

### Retrieving the list of volumes

```
# The API:
/api/storage/volumes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "2d1167cc-c3f2-495a-a23f-8f50b071b9b8",
    "name": "vsdata_root",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/2d1167cc-c3f2-495a-a23f-
8f50b071b9b8"
      }
    }
  },
  {
    "uuid": "3969be7e-78b4-4b4c-82a4-fa86331f03df",
    "name": "vsfg_root",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/3969be7e-78b4-4b4c-82a4-
fa86331f03df"
      }
    }
  },
  {
    "uuid": "59c03ac5-e708-4ce8-a676-278dc249fda2",
    "name": "svm_root",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/59c03ac5-e708-4ce8-a676-
278dc249fda2"
      }
    }
```

```
  },
  {
    "uuid": "6802635b-8036-11e8-aae5-0050569503ac",
    "name": "fgvol",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/6802635b-8036-11e8-aae5-
0050569503ac"
      }
    }
  },
  {
    "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
    "name": "datavol",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-
e3295a7e9057"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes"
  }
}
}
```

**Retrieving the attributes of a volume**

The GET request is used to retrieve the attributes of a volume.

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -H "accept: application/hal+json"

# The response:
{
"uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
"comment": "This is a data volume",
"create_time": "2018-07-05T14:56:44+05:30",
```

```
"language": "en_us",
"name": "datavol",
"size": 20971520,
"state": "online",
"style": "flexvol",
"tiering_policy": "auto",
"type": "rw",
"aggregates": [
  {
    "name": "data",
    "uuid": "aa742322-36bc-4d98-bbc4-0a827534c035",
    "_links": {
      "self": {
        "href": "/api/cluster/aggregates/data"
      }
    }
  }
],
"efficiency": {
  "ratio": 1
},
"encryption": {
  "enabled": false,
  "state": "none",
  "key_id": "",
  "type" : "none"
},
"error_state": {
  "has_bad_blocks": false,
  "is_inconsistent": false
},
"files": {
  "maximum": 566,
  "used": 96
},
"nas": {
  "gid": 2468,
  "security_style": "unix",
  "uid": 1357,
  "unix_permissions": 4755,
  "export_policy": {
    "name": "default",
    "id": 8589934593
  },
  "junction_parent": {
    "name": "vol1",
```

```
    "uuid": "a2564f80-25fb-41e8-9b49-44de2600991f",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/a2564f80-25fb-41e8-9b49-
44de2600991f"
      }
    }
  }
},
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "status": "ok",
  "duration": "PT15S",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "cloud": {
    "timestamp": "2019-04-09T05:50:15Z",
    "status": "ok",
    "duration": "PT15S",
    "iops" : {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "latency": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
```

```
      }
    },
    "flexcache": {
      "timestamp": "2019-04-09T05:50:15Z",
      "status": "ok",
      "duration": "PT1D",
      "cache_miss_percent": 0,
      "bandwidth_savings": 0
    }
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "latency_raw": {
      "other": 38298,
      "total": 38298,
      "read": 0,
      "write": 0
    },
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 3,
      "total": 3
    },
    "throughput_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "cloud": {
      "timestamp": "2019-04-09T05:50:42Z",
      "status": "ok",
      "iops_raw" : {
        "read": 0,
        "write": 0,
        "other": 0,
        "total": 0
      },
      "latency_raw": {
        "read": 0,
        "write": 0,
        "other": 0,
        "total": 0
      }
```

```
    },
    "flexcache_raw": {
      "timestamp": "2019-04-09T05:50:15Z",
      "status": "ok",
      "cache_miss_blocks": 0,
      "client_requested_blocks": 0
    }
  },
  "qos": {
    "policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "228454af-5a8b-11e9-bd5b-005056ac6f1f",
    "name": "pg1"
    }
  },
  "snaplock": {
    "append_mode_enabled": false,
    "autocommit_period": "none",
    "compliance_clock_time": "2019-05-24T10:59:00+05:30",
    "expiry_time": "2038-01-19T08:44:28+05:30",
    "is_audit_log": false,
    "litigation_count": 0,
    "privileged_delete": "disabled",
    "type": "enterprise",
    "retention": {
      "default": "P0Y",
      "minimum": "P0Y",
      "maximum": "P30Y"
    }
  },
  "snapshot_policy": {
    "name": "default"
  },
  "svm": {
    "name": "vsdata",
    "uuid": "d61b69f5-7458-11e8-ad3f-0050569503ac"
  }
  }
```


**Retrieving the quota state of a FlexVol or a FlexGroup volume**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb20da45-4f6b-11e9-
9a71-005056a7f717/?fields=quota.state" -H "accept: application/hal+json"

# The response:
{
"uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
"name": "fv",
"quota": {
  "state": "on"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717/"
  }
}
}
```

**Retrieving the constituents of a FlexGroup volume**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes?flexgroup.uuid=fd87d06f-8876-11ec-94a3-
005056a7484f&is_constituent=true" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "fd877f7c-8876-11ec-94a3-005056a7484f",
    "name": "fg__0001",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/fd877f7c-8876-11ec-94a3-
005056a7484f?is_constituent=true"
```

```
      }
    }
  },
  {
    "uuid": "fea631d6-8876-11ec-94a3-005056a7484f",
    "name": "fg__0002",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/fea631d6-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  },
  {
    "uuid": "ff38a34e-8876-11ec-94a3-005056a7484f",
    "name": "fg__0003",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ff38a34e-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  },
  {
    "uuid": "ffdbbd1f-8876-11ec-94a3-005056a7484f",
    "name": "fg__0004",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ffdbbd1f-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
```

```
    "href": "/api/storage/volumes?flexgroup.uuid=fd87d06f-8876-11ec-94a3-
005056a7484f&is_constituent=true"
  }
}
}
```


**Retrieving the efficiency attributes of volume**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/5f098ebc-32c8-11eb-
8dde-005056ace228/?fields=efficiency" -H "accept: application/hal+json"

# The response:
{
"uuid": "5f098ebc-32c8-11eb-8dde-005056ace228",
"name": "vol1",
"efficiency": {
  "compression": "both",
  "dedupe": "background",
  "cross_volume_dedupe": "none",
  "compaction": "none",
  "schedule": "sun-sat@0",
  "svm": "vs0",
  "state": "enabled",
  "status": "idle",
  "type": "regular",
  "progress": "Idle for 00:10:37",
  "last_op_begin_timestamp": "Mon Nov 30 00:00:02 2020",
  "last_op_end_timestamp": "Mon Nov 30 00:00:03 2020",
  "last_op_state": "Success",
  "last_op_size": 0,
  "addr": "/vol/vol1",
  "policy": {
    "name": "-"
  },
  "ratio": 1
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/5f098ebc-32c8-11eb-8dde-005056ace228"
   }
}
}
```

#### Updating the attributes of a volume

#### Examples

### Updating the attributes of a volume

The PATCH request is used to update the attributes of a volume.

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{ "size": 26214400, "nas": {"security_style":
"mixed"}, "comment": "This is a data volume" }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 31 Jul 2018 09:36:43 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "3c5be5a6-94a5-11e8-8ca3-00505695c11b",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/3c5be5a6-94a5-11e8-8ca3-00505695c11b"
    }
  }
}
}
```

**Updating the attributes of a FlexClone using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"clone":{"split_initiated":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "8e01747f-38e5-11e9-8a3a-00a09843994b",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8e01747f-38e5-11e9-8a3a-00a09843994b"
    }
  }
}
}
```

**Stopping a volume clone split operation on a FlexClone using PATCH.**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"clone":{"split_initiated":"false"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 03 Nov 2021 15:10:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "1f8b3673-3cb8-11ec-b89e-005056bb2cb5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/1f8b3673-3cb8-11ec-b89e-005056bb2cb5"
    }
  }
}
}
```

**Enabling quotas for a FlexVol or a FlexGroup volume using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "d2fe7299-57d0-11e9-a2dc-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/d2fe7299-57d0-11e9-a2dc-005056a7f717"
    }
  }
}
}
```

**Disabling quotas for a FlexVol or a FlexGroup volume using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"false"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "0c8f6bea-57d1-11e9-a2dc-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/0c8f6bea-57d1-11e9-a2dc-005056a7f717"
    }
  }
}
}
```

**Starting non-disruptive volume capacity rebalancing for a FlexGroup volume using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"state":"starting",
"max_runtime":"PT6H"}}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "d2fe7299-57d0-11e9-a2dc-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/d2fe7299-57d0-11e9-a2dc-005056a7f717"
    }
  }
}
}
```

**Starting a scheduled non-disruptive volume capacity rebalancing for a FlexGroup volume using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"state":"starting",
"start_time":"2022-12-21T15:30:00-05:00"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 21 Dec 2022 15:25:00 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "53b2a954-80d4-11ed-81d2-005056ac8ed0",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/53b2a954-80d4-11ed-81d2-005056ac8ed0"
    }
  }
}
}
```

**Stopping non-disruptive volume capacity rebalancing OR scheduled rebalancing for a FlexGroup volume using PATCH. This works for scheduled or on-going rebalancing.**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"state":"stopping"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "0c8f6bea-57d1-11e9-a2dc-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/0c8f6bea-57d1-11e9-a2dc-005056a7f717"
    }
  }
}
}
```

**Modifying non-disruptive volume capacity rebalancing configurations for a FlexGroup volume**

The following example shows how to use a PATCH request to modify non-disruptive volume capacity rebalancing configurations for a FlexGroup volume:

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"start_time":"2023-03-18T15:30:00-
05:00","max_threshold":20, "min_threshold":5, "max_file_moves":15,
"min_file_size":"100MB", "exclude_snapshots":"false",
"max_runtime":"PT6H"}}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 16 Mar 2023 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "0c8f6bea-57d1-11e9-a2dc-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/0c8f6bea-57d1-11e9-a2dc-005056a7f717"
    }
  }
}
}
```

#### Retrieving non-disruptive volume capacity rebalancing engine runtime information for a FlexGroup volume

The following example shows how to use a GET request to retrieve non-disruptive volume capacity rebalancing engine runtime information for a FlexGroup volume:

```
# The API:
/api/storage/volumes?is_constituent=true

# The call:
curl -X GET  "https://<mgmt-
ip>/api/storage/volumes?fields=rebalancing.engine&is_constituent=true&flex
group.uuid=d0c3359c-5448-4a9b-a077-e3295a7e9057" -H "accept:
application/hal+json"

# The response:
```

```
{
"records": [
  {
    "uuid": "2b32fdf1-b916-11ec-b103-005056a79638",
    "name": "fg__0001",
    "flexgroup": {
      "uuid": "2b3323db-b916-11ec-b103-005056a79638"
    },
    "rebalancing": {
      "engine": {
        "scanner": {
          "files_scanned": 3522915,
          "files_skipped": {
            "too_small": 3812,
            "too_large": 199,
            "fast_truncate": 22,
            "in_snapshot": 77499,
            "efficiency_blocks": 1823,
            "efficiency_percent": 355,
            "incompatible": 9377,
            "metadata": 85449,
            "remote_cache": 1912,
            "write_fenced": 28,
            "on_demand_destination": 87,
            "footprint_invalid": 12,
            "other": 336
          },
          "blocks_scanned": 1542675000,
          "blocks_skipped": {
            "too_small": 8744000,
            "too_large": 865000,
            "fast_truncate": 54000,
            "in_snapshot": 7749000,
            "efficiency_blocks": 1472000,
            "efficiency_percent": 366000,
            "incompatible": 2287000,
            "metadata": 85673000,
            "remote_cache": 9914000,
            "write_fenced": 19000,
            "on_demand_destination": 66000,
            "footprint_invalid": 98000,
            "other":187000
          }
        },
        "movement": {
          "file_moves_started": 9833,
```

```
          "most_recent_start_time": "2022-02-15T12:56:07-05:00",
          "last_error": {
            "time": "2022-02-15T09:09:27-05:00",
            "file_id": 88,
            "destination": 1089,
            "code": 60
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/2b32fdf1-b916-11ec-b103-
005056a79638?is_constituent=true"
      }
    }
  },
  {
    "uuid": "2cc5da55-b916-11ec-b103-005056a79638",
    "name": "fg__0002",
    "flexgroup": {
      "uuid": "2b3323db-b916-11ec-b103-005056a79638"
    },
    "rebalancing": {
      "engine": {
        "scanner": {
          "files_scanned": 3522915,
          "files_skipped": {
            "too_small": 3812,
            "too_large": 188,
            "fast_truncate": 25,
            "in_snapshot": 77499,
            "efficiency_blocks": 1823,
            "efficiency_percent": 355,
            "incompatible": 9377,
            "metadata": 85449,
            "remote_cache": 1912,
            "write_fenced": 28,
            "on_demand_destination": 87,
            "footprint_invalid": 12,
            "other": 336
          },
          "blocks_scanned": 1542675000,
          "blocks_skipped": {
            "too_small": 8744000,
            "too_large": 865000,
```

```
            "fast_truncate": 54000,
            "in_snapshot": 7749000,
            "efficiency_blocks": 1472000,
            "efficiency_percent": 366000,
            "incompatible": 2287000,
            "metadata": 85673000,
            "remote_cache": 9914000,
            "write_fenced": 19000,
            "on_demand_destination": 66000,
            "footprint_invalid": 98000,
            "other":187000
          }
        },
        "movement": {
          "file_moves_started": 9833,
          "most_recent_start_time": "2022-02-15T12:56:07-05:00",
          "last_error": {
            "time": "2022-02-15T08:09:27-05:00",
            "file_id": 88,
            "destination": 1089,
            "code": 60
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/2cc5da55-b916-11ec-b103-
005056a79638?is_constituent=true"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/storage/volumes?fields=rebalancing.engine&is_constituent=true&flexgr
oup.uuid=d0c3359c-5448-4a9b-a077-e3295a7e9057"
  }
}
}
```


#### Add tiering object tags for a FlexVol volume

The following example shows how to use a PATCH request to add tiering object tags for a FlexVol volume:

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": [ "key1=val1", "key2=val2"
]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 11 Feb 2020 19:29:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "d05012de-4d04-11ea-836b-005056bb6f9d",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/d05012de-4d04-11ea-836b-005056bb6f9d"
    }
  }
}
}
```

**Remove tiering object tags for a FlexVol using PATCH**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": []}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 Jan 2020 22:28:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "ca234df1-3ef8-11ea-9a56-005056bb69a1",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/ca234df1-3ef8-11ea-9a56-005056bb69a1"
    }
  }
}
}
```

#### Deleting a volume

#### Example

### Deleting a volume

The DELETE request is used to delete a volume.

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X DELETE  "https://<mgmt-ip>/api/storage/volumes/{uuid} " -H
"accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
content-length: 189
content-type: application/json
date: Wed, 01 Aug 2018 09:40:36 GMT
keep-alive: timeout=5, max=100
server: libzapid-httpd
{
"job": {
  "uuid": "f1aa3eb8-956e-11e8-86bf-0050568e2249",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f1aa3eb8-956e-11e8-86bf-0050568e2249"
    }
  }
}
}
```

**Deleting a volume and bypassing the recovery queue**

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X DELETE  "https://<mgmt-ip>/api/storage/volumes/{uuid}?force=true"
-H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 09 Feb 2022 09:59:55 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "08757020-898f-11ec-b367-005056bb7353",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/08757020-898f-11ec-b367-005056bb7353"
    }
  }
}
}
```

### Retrieve volumes

GET `/storage/volumes`

**Introduced In:** 9.6

Retrieves volumes.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `is_svm_root`
- `aggressive_readahead_mode`
- `analytics.by_accessed_time.bytes_used`

- analytics.by_modified_time.bytes_used
- analytics.bytes_used
- analytics.file_count
- analytics.histogram_by_time_labels
- analytics.incomplete_data
- analytics.report_time
- analytics.subdir_count
- analytics.supported
- analytics.unsupported_reason
- anti_ransomware.*
- application.*
- encryption.*
- queue_for_encryption
- convert_unicode
- clone.parent_snapshot.name
- clone.parent_snapshot.uuid
- clone.parent_svm.name
- clone.parent_svm.uuid
- clone.parent_volume.name
- clone.parent_volume.uuid
- clone.split_complete_percent
- clone.split_estimate
- clone.split_initiated
- efficiency.*
- error_state.*
- files.*
- max_dir_size
- nas.export_policy.id
- nas.gid
- nas.path
- nas.security_style
- nas.uid
- nas.unix_permissions
- nas.junction_parent.name

- nas.junction_parent.uuid
- snaplock.*
- restore_to.*
- quota.*
- qos.*
- flexcache_endpoint_type
- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint
- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.dedupe_metafiles_footprint
- space.dedupe_metafiles_temporary_footprint
- space.delayed_free_footprint
- space.file_operation_metadata
- space.snapmirror_destination_footprint
- space.volume_guarantee_footprint
- space.cross_volume_dedupe_metafiles_footprint
- space.cross_volume_dedupe_metafiles_temporary_footprint
- space.snapshot_reserve_unusable
- space.snapshot_spill
- space.user_data
- space.logical_space.*
- space.snapshot.*
- space.used_by_afs
- space.afs_total
- space.available_percent
- space.full_threshold_percent
- space.nearly_full_threshold_percent
- space.overwrite_reserve
- space.overwrite_reserve_used

- space.size_available_for_snapshots
- space.percent_used
- space.fractional_reserve
- space.block_storage_inactive_user_data_percent
- space.physical_used
- space.physical_used_percent
- space.expected_available
- space.filesystem_size
- space.filesystem_size_fixed
- guarantee.*
- autosize.*
- movement.*
- statistics.*
- constituents.name
- constituents.node.name
- constituents.node.uuid
- constituents.space.size
- constituents.space.available
- constituents.space.used
- constituents.space.available_percent
- constituents.space.used_percent
- constituents.space.block_storage_inactive_user_data
- constituents.space.capacity_tier_footprint
- constituents.space.performance_tier_footprint
- constituents.space.local_tier_footprint
- constituents.space.footprint
- constituents.space.over_provisioned
- constituents.space.metadata
- constituents.space.total_footprint
- constituents.space.logical_space.reporting
- constituents.space.logical_space.enforcement
- constituents.space.logical_space.used_by_afs
- constituents.space.logical_space.available
- constituents.space.snapshot.used

- constituents.space.snapshot.reserve_percent
- constituents.space.snapshot.autodelete_enabled
- constituents.space.large_size_enabled
- constituents.space.max_size
- constituents.space.total_metadata
- constituents.space.total_metadata_footprint
- constituents.aggregates.name
- constituents.aggregates.uuid
- constituents.movement.destination_aggregate.name
- constituents.movement.destination_aggregate.uuid
- constituents.movement.state
- constituents.movement.percent_complete
- constituents.movement.cutover_window
- constituents.movement.tiering_policy
- asynchronous_directory_delete.*
- rebalancing.*
- metric.*
- cloud_write_enabled

#### Related ONTAP commands

- volume show
- volume clone show
- volume efficiency show
- volume encryption show
- volume flexcache show
- volume flexgroup show
- volume move show
- volume quota show
- volume show-space
- volume snaplock show
- volume rebalance show
- security anti-ransomware volume show
- security anti-ransomware volume space show
- volume file async-delete client show

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| is_constituent | boolean | query | False | When set to false, only FlexVol and FlexGroup volumes are returned. When set to true, only FlexGroup constituent volumes are returned. Default for GET calls is false.<br>• Default value:<br>• Introduced in: 9.10 |
| snapmirror.is_protected | boolean | query | False | Filter by snapmirror.is_protected<br>• Introduced in: 9.7 |
| snapmirror.destinations.is_cloud | boolean | query | False | Filter by snapmirror.destinations.is_cloud<br>• Introduced in: 9.9 |
| snapmirror.destinations.is_ontap | boolean | query | False | Filter by snapmirror.destinations.is_ontap<br>• Introduced in: 9.9 |
| anti_ransomware.clear_suspect.phase | string | query | False | Filter by anti_ransomware.clear_suspect.phase<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.clear_suspect.start_time | string | query | False | Filter by anti_ransomware.clear_suspect.start_time<br><br>• Introduced in: 9.18 |
| anti_ransomware.attack_probability | string | query | False | Filter by anti_ransomware.attack_probability<br><br>• Introduced in: 9.10 |
| anti_ransomware.event_log.is_enabled_on_new_file_extension_seen | boolean | query | False | Filter by anti_ransomware.event_log.is_enabled_on_new_file_extension_seen<br><br>• Introduced in: 9.14 |
| anti_ransomware.event_log.is_enabled_on_snapshot_copy_creation | boolean | query | False | Filter by anti_ransomware.event_log.is_enabled_on_snapshot_copy_creation<br><br>• Introduced in: 9.14 |
| anti_ransomware.typical_usage.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.typical_usage.file_create_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.typical_usage.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.typical_usage.file_rename_peak_rate_per_minute<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.typical_usage.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.typical_usage.high_entropy_data_write_peak_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.typical_usage.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.typical_usage.file_delete_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_reports.time | string | query | False | Filter by anti_ransomware.attack_reports.time<br><br>• Introduced in: 9.10 |
| anti_ransomware.block_device_detection_start_time | string | query | False | Filter by anti_ransomware.block_device_detection_start_time<br><br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.attack_detection_parameters.relaxing_popular_file_extensions | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.relaxing_popular_file_extensions<br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.based_on_file_delete_op_rate | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.based_on_file_delete_op_rate<br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.file_delete_op_rate_surge_notify_percent | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.file_delete_op_rate_surge_notify_percent<br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.file_rename_op_rate_surge_notify_percent | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.file_rename_op_rate_surge_notify_percent<br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.block_device_auto_learned_encryption_threshold | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.block_device_auto_learned_encryption_threshold<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| anti_ransomware.attack_detection_parameters.based_on_file_rename_op_rate | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.based_on_file_rename_op_rate<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.high_entropy_data_surge_notify_percent | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.high_entropy_data_surge_notify_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.file_create_op_rate_surge_notify_percent | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.file_create_op_rate_surge_notify_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.based_on_file_create_op_rate | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.based_on_file_create_op_rate<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_duration_in_hours | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_duration_in_hours<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.attack_detection_parameters.based_on_never_seen_before_file_extension | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.based_on_never_seen_before_file_extension<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.based_on_high_entropy_data_rate | boolean | query | False | Filter by anti_ransomware.attack_detection_parameters.based_on_high_entropy_data_rate<br><br>• Introduced in: 9.16 |
| anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_count_notify_threshold | integer | query | False | Filter by anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_count_notify_threshold<br><br>• Introduced in: 9.16 |
| anti_ransomware.suspect_files.count | integer | query | False | Filter by anti_ransomware.suspect_files.count<br><br>• Introduced in: 9.10 |
| anti_ransomware.suspect_files.entropy | string | query | False | Filter by anti_ransomware.suspect_files.entropy<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| anti_ransomware.suspect_files.format | string | query | False | Filter by anti_ransomware.suspect_files.format<br><br>• Introduced in: 9.10 |
| anti_ransomware.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.surge_usage.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.surge_usage.file_delete_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.surge_usage.time | string | query | False | Filter by anti_ransomware.surge_usage.time<br><br>• Introduced in: 9.16 |
| anti_ransomware.surge_usage.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.surge_usage.high_entropy_data_write_peak_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.surge_usage.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.surge_usage.file_create_peak_rate_per_minute<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.surge_usage.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.surge_usage.file_rename_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.block_device_detection_state | string | query | False | Filter by anti_ransomware.block_device_detection_state<br><br>• Introduced in: 9.17 |
| anti_ransomware.workload.typical_usage.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.typical_usage.file_delete_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.typical_usage.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.typical_usage.file_rename_peak_rate_per_minute<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.workload.typical_usage.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.typical_usage.file_create_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.newly_observed_file_extensions.count | integer | query | False | Filter by anti_ransomware.workload.newly_observed_file_extensions.count<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.newly_observed_file_extensions.name | string | query | False | Filter by anti_ransomware.workload.newly_observed_file_extensions.name<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.file_extension_types_count | integer | query | False | Filter by anti_ransomware.workload.file_extension_types_count<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_percent<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_usage.file_rename_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_usage.file_create_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_usage.file_delete_peak_rate_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.newly_observed_file_extensions.name | string | query | False | Filter by anti_ransomware.workload.surge_usage.newly_observed_file_extensions.name<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.workload.surge_usage.newly_observed_file_extensions.count | integer | query | False | Filter by anti_ransomware.workload.surge_usage.newly_observed_file_extensions.count<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_usage.time | string | query | False | Filter by anti_ransomware.workload.surge_usage.time<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_statistics.time | string | query | False | Filter by anti_ransomware.workload.surge_statistics.time<br><br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_rate_kb_per_minute<br><br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.workload.surge_statistics.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_statistics.file_delete_peak_rate_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_statistics.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_statistics.file_create_peak_rate_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_statistics.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.surge_statistics.file_rename_peak_rate_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_percent<br>• Introduced in: 9.16 |
| anti_ransomware.workload.historical_statistics.file_delete_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.historical_statistics.file_delete_peak_rate_per_minute<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_rate_kb_per_minute | integer | query | False | Filter by anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_rate_kb_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_percent | integer | query | False | Filter by anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_percent<br>• Introduced in: 9.16 |
| anti_ransomware.workload.historical_statistics.file_rename_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.historical_statistics.file_rename_peak_rate_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.historical_statistics.file_create_peak_rate_per_minute | integer | query | False | Filter by anti_ransomware.workload.historical_statistics.file_create_peak_rate_per_minute<br>• Introduced in: 9.16 |
| anti_ransomware.workload.file_extensions_observed | string | query | False | Filter by anti_ransomware.workload.file_extensions_observed<br>• Introduced in: 9.16 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| anti_ransomware.space.used | integer | query | False | Filter by anti_ransomware.space.used<br><br>• Introduced in: 9.10 |
| anti_ransomware.space.used_by_snapshots | integer | query | False | Filter by anti_ransomware.space.used_by_snapshots<br><br>• Introduced in: 9.10 |
| anti_ransomware.space.used_by_logs | integer | query | False | Filter by anti_ransomware.space.used_by_logs<br><br>• Introduced in: 9.10 |
| anti_ransomware.space.snapshot_count | integer | query | False | Filter by anti_ransomware.space.snapshot_count<br><br>• Introduced in: 9.10 |
| anti_ransomware.surge_as_normal | boolean | query | False | Filter by anti_ransomware.surge_as_normal<br><br>• Introduced in: 9.11 |
| anti_ransomware.attack_detected_by | string | query | False | Filter by anti_ransomware.attack_detected_by<br><br>• Introduced in: 9.17 |
| anti_ransomware.dry_run_start_time | string | query | False | Filter by anti_ransomware.dry_run_start_time<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| anti_ransomware.state | string | query | False | Filter by anti_ransomware.state<br>• Introduced in: 9.10 |
| anti_ransomware.update_baseline_from_surge | boolean | query | False | Filter by anti_ransomware.update_baseline_from_surge<br>• Introduced in: 9.15 |
| is_dir_index_transfer_enabled | boolean | query | False | Filter by is_dir_index_transfer_enabled<br>• Introduced in: 9.17 |
| granular_data_mode | string | query | False | Filter by granular_data_mode<br>• Introduced in: 9.16 |
| is_object_store | boolean | query | False | Filter by is_object_store<br>• Introduced in: 9.8 |
| comment | string | query | False | Filter by comment<br>• maxLength: 1023<br>• minLength: 0 |
| clone.qtree_name | string | query | False | Filter by clone.qtree_name<br>• Introduced in: 9.18 |
| clone.split_estimate | integer | query | False | Filter by clone.split_estimate |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| clone.is_flexclone | boolean | query | False | Filter by clone.is_flexclone |
| clone.lun_name | string | query | False | Filter by clone.lun_name<br>• Introduced in: 9.18 |
| clone.split_initiated | boolean | query | False | Filter by clone.split_initiated |
| clone.parent_svm.name | string | query | False | Filter by clone.parent_svm.name |
| clone.parent_svm.uuid | string | query | False | Filter by clone.parent_svm.uuid |
| clone.split_complete_percent | integer | query | False | Filter by clone.split_complete_percent |
| clone.parent_volume.name | string | query | False | Filter by clone.parent_volume.name |
| clone.parent_volume.uuid | string | query | False | Filter by clone.parent_volume.uuid |
| clone.parent_snapshot.name | string | query | False | Filter by clone.parent_snapshot.name |
| clone.parent_snapshot.uuid | string | query | False | Filter by clone.parent_snapshot.uuid |
| clone.inherited_savings | integer | query | False | Filter by clone.inherited_savings<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| clone.has_flexclone | boolean | query | False | Filter by clone.has_flexclone<br>• Introduced in: 9.16 |
| clone.inherited_physical_used | integer | query | False | Filter by clone.inherited_physical_used<br>• Introduced in: 9.12 |
| language | string | query | False | Filter by language |
| files.used | integer | query | False | Filter by files.used |
| files.maximum | integer | query | False | Filter by files.maximum |
| files.inodefile_capacity | integer | query | False | Filter by files.inodefile_capacity<br>• Introduced in: 9.17 |
| flexgroup.uuid | string | query | False | Filter by flexgroup.uuid<br>• Introduced in: 9.10 |
| flexgroup.name | string | query | False | Filter by flexgroup.name<br>• Introduced in: 9.10<br>• maxLength: 203<br>• minLength: 1 |
| size | integer | query | False | Filter by size |
| style | string | query | False | Filter by style |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| nodes.name | string | query | False | Filter by nodes.name<br>• Introduced in: 9.17 |
| nodes.uuid | string | query | False | Filter by nodes.uuid<br>• Introduced in: 9.17 |
| application.uuid | string | query | False | Filter by application.uuid |
| application.name | string | query | False | Filter by application.name |
| aggressive_readahead_mode | string | query | False | Filter by aggressive_readahead_mode<br>• Introduced in: 9.13 |
| create_time | string | query | False | Filter by create_time |
| consistency_group.name | string | query | False | Filter by consistency_group.name<br>• Introduced in: 9.7 |
| consistency_group.uuid | string | query | False | Filter by consistency_group.uuid<br>• Introduced in: 9.10 |
| access_time_enabled | boolean | query | False | Filter by access_time_enabled<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| granular_data | boolean | query | False | Filter by granular_data<br>• Introduced in: 9.11 |
| is_svm_root | boolean | query | False | Filter by is_svm_root<br>• Introduced in: 9.7 |
| is_s3_arbitrary_part_size_enabled | boolean | query | False | Filter by is_s3_arbitrary_part_size_enabled<br>• Introduced in: 9.18 |
| has_large_dir | boolean | query | False | Filter by has_large_dir<br>• Introduced in: 9.18 |
| quota.state | string | query | False | Filter by quota.state |
| queue_for_encryption | boolean | query | False | Filter by queue_for_encryption<br>• Introduced in: 9.8 |
| has_dir_index_public | boolean | query | False | Filter by has_dir_index_public<br>• Introduced in: 9.17 |
| name | string | query | False | Filter by name<br>• maxLength: 203<br>• minLength: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| asynchronous_directory_delete.trash_bin | string | query | False | Filter by asynchronous_directory_delete.trash_bin<br><br>• Introduced in: 9.11 |
| asynchronous_directory_delete.enabled | boolean | query | False | Filter by asynchronous_directory_delete.enabled<br><br>• Introduced in: 9.11 |
| rebalancing.target_used | integer | query | False | Filter by rebalancing.target_used<br><br>• Introduced in: 9.11 |
| rebalancing.imbalance_size | integer | query | False | Filter by rebalancing.imbalance_size<br><br>• Introduced in: 9.11 |
| rebalancing.min_file_size | integer | query | False | Filter by rebalancing.min_file_size<br><br>• Introduced in: 9.11 |
| rebalancing.used_for_imbalance | integer | query | False | Filter by rebalancing.used_for_imbalance<br><br>• Introduced in: 9.12 |
| rebalancing.notices.arguments.message | string | query | False | Filter by rebalancing.notices.arguments.message<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.notices.arguments.code | string | query | False | Filter by rebalancing.notices.arguments.code<br><br>• Introduced in: 9.12 |
| rebalancing.notices.code | string | query | False | Filter by rebalancing.notices.code<br><br>• Introduced in: 9.12 |
| rebalancing.notices.message | string | query | False | Filter by rebalancing.notices.message<br><br>• Introduced in: 9.12 |
| rebalancing.runtime | string | query | False | Filter by rebalancing.runtime<br><br>• Introduced in: 9.11 |
| rebalancing.stop_time | string | query | False | Filter by rebalancing.stop_time<br><br>• Introduced in: 9.11 |
| rebalancing.state | string | query | False | Filter by rebalancing.state<br><br>• Introduced in: 9.11 |
| rebalancing.start_time | string | query | False | Filter by rebalancing.start_time<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.data_moved | integer | query | False | Filter by rebalancing.data_moved<br><br>• Introduced in: 9.11 |
| rebalancing.min_threshold | integer | query | False | Filter by rebalancing.min_threshold<br><br>• Introduced in: 9.11 |
| rebalancing.max_file_moves | integer | query | False | Filter by rebalancing.max_file_moves<br><br>• Introduced in: 9.11 |
| rebalancing.max_threshold | integer | query | False | Filter by rebalancing.max_threshold<br><br>• Introduced in: 9.11 |
| rebalancing.max_constituent_imbalance_percent | integer | query | False | Filter by rebalancing.max_constituent_imbalance_percent<br><br>• Introduced in: 9.11 |
| rebalancing.imbalance_percent | integer | query | False | Filter by rebalancing.imbalance_percent<br><br>• Introduced in: 9.11 |
| rebalancing.exclude_snapshots | boolean | query | False | Filter by rebalancing.exclude_snapshots<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.engine.scanner.blocks_skipped.remote_cache | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.remote_cache<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.write_fenced | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.write_fenced<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.in_snapshot | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.in_snapshot<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.metadata | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.metadata<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.fast_truncate | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.fast_truncate<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.other | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.other<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| rebalancing.engine.scanner.blocks_skipped.efficiency_blocks | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_blocks<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.too_large | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.too_large<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.too_small | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.too_small<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.efficiency_percent | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_percent<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.on_demand_destination | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.on_demand_destination<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_skipped.footprint_invalid | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.footprint_invalid<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.engine.scanner.blocks_skipped.incompatible | integer | query | False | Filter by rebalancing.engine.scanner.blocks_skipped.incompatible<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.blocks_scanned | integer | query | False | Filter by rebalancing.engine.scanner.blocks_scanned<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.too_large | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.too_large<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.too_small | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.too_small<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.incompatible | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.incompatible<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.efficiency_percent | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.efficiency_percent<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.engine.scanner.files_skipped.footprint_invalid | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.footprint_invalid<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.on_demand_destination | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.on_demand_destination<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.write_fenced | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.write_fenced<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.remote_cache | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.remote_cache<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.other | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.other<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.efficiency_blocks | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.efficiency_blocks<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.engine.scanner.files_skipped.in_snapshot | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.in_snapshot<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.metadata | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.metadata<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_skipped.fast_truncate | integer | query | False | Filter by rebalancing.engine.scanner.files_skipped.fast_truncate<br><br>• Introduced in: 9.12 |
| rebalancing.engine.scanner.files_scanned | integer | query | False | Filter by rebalancing.engine.scanner.files_scanned<br><br>• Introduced in: 9.12 |
| rebalancing.engine.movement.last_error.destination | integer | query | False | Filter by rebalancing.engine.movement.last_error.destination<br><br>• Introduced in: 9.12 |
| rebalancing.engine.movement.last_error.file_id | integer | query | False | Filter by rebalancing.engine.movement.last_error.file_id<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| rebalancing.engine.movement.last_error.time | string | query | False | Filter by rebalancing.engine.movement.last_error.time<br><br>• Introduced in: 9.12 |
| rebalancing.engine.movement.last_error.code | integer | query | False | Filter by rebalancing.engine.movement.last_error.code<br><br>• Introduced in: 9.12 |
| rebalancing.engine.movement.file_moves_started | integer | query | False | Filter by rebalancing.engine.movement.file_moves_started<br><br>• Introduced in: 9.12 |
| rebalancing.engine.movement.most_recent_start_time | string | query | False | Filter by rebalancing.engine.movement.most_recent_start_time<br><br>• Introduced in: 9.12 |
| rebalancing.max_runtime | string | query | False | Filter by rebalancing.max_runtime<br><br>• Introduced in: 9.11 |
| error_state.has_bad_blocks | boolean | query | False | Filter by error_state.has_bad_blocks |
| error_state.is_inconsistent | boolean | query | False | Filter by error_state.is_inconsistent |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| snapshot_policy.uuid | string | query | False | Filter by snapshot_policy.uuid |
| snapshot_policy.name | string | query | False | Filter by snapshot_policy.name |
| constituents.name | string | query | False | Filter by constituents.name<br>• Introduced in: 9.9 |
| constituents.movement.state | string | query | False | Filter by constituents.movement.state<br>• Introduced in: 9.9 |
| constituents.movement.cutover_window | integer | query | False | Filter by constituents.movement.cutover_window<br>• Introduced in: 9.9 |
| constituents.movement.percent_complete | integer | query | False | Filter by constituents.movement.percent_complete<br>• Introduced in: 9.9 |
| constituents.movement.destination_aggregate.name | string | query | False | Filter by constituents.movement.destination_aggregate.name<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| constituents.movement.destination_aggregate.uuid | string | query | False | Filter by constituents.movement.destination_aggregate.uuid<br><br>• Introduced in: 9.9 |
| constituents.aggregates.uuid | string | query | False | Filter by constituents.aggregates.uuid<br><br>• Introduced in: 9.9 |
| constituents.aggregates.name | string | query | False | Filter by constituents.aggregates.name<br><br>• Introduced in: 9.9 |
| constituents.space.total_metadata_footprint | integer | query | False | Filter by constituents.space.total_metadata_footprint<br><br>• Introduced in: 9.15 |
| constituents.space.used_percent | integer | query | False | Filter by constituents.space.used_percent<br><br>• Introduced in: 9.10 |
| constituents.space.performance_tier_footprint | integer | query | False | Filter by constituents.space.performance_tier_footprint<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| constituents.space.logical_space.reporting | boolean | query | False | Filter by constituents.space.logical_space.reporting<br><br>• Introduced in: 9.9 |
| constituents.space.logical_space.used_by_afs | integer | query | False | Filter by constituents.space.logical_space.used_by_afs<br><br>• Introduced in: 9.9 |
| constituents.space.logical_space.available | integer | query | False | Filter by constituents.space.logical_space.available<br><br>• Introduced in: 9.9 |
| constituents.space.logical_space.enforcement | boolean | query | False | Filter by constituents.space.logical_space.enforcement<br><br>• Introduced in: 9.9 |
| constituents.space.total_footprint | integer | query | False | Filter by constituents.space.total_footprint<br><br>• Introduced in: 9.9 |
| constituents.space.used_by_afs | integer | query | False | Filter by constituents.space.used_by_afs<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| constituents.space.used | integer | query | False | Filter by constituents.space.used<br><br>• Introduced in: 9.9 |
| constituents.space.footprint | integer | query | False | Filter by constituents.space.footprint<br><br>• Introduced in: 9.9 |
| constituents.space.over_provisioned | integer | query | False | Filter by constituents.space.over_provisioned<br><br>• Introduced in: 9.9 |
| constituents.space.large_size_enabled | boolean | query | False | Filter by constituents.space.large_size_enabled<br><br>• Introduced in: 9.12 |
| constituents.space.snapshot.used | integer | query | False | Filter by constituents.space.snapshot.used<br><br>• Introduced in: 9.9 |
| constituents.space.snapshot.reserve_percent | integer | query | False | Filter by constituents.space.snapshot.reserve_percent<br><br>• Introduced in: 9.9 |
| constituents.space.available | integer | query | False | Filter by constituents.space.available<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| constituents.space.size | integer | query | False | Filter by constituents.space.size<br><br>• Introduced in: 9.9 |
| constituents.space.metadata | integer | query | False | Filter by constituents.space.metadata<br><br>• Introduced in: 9.9 |
| constituents.space.block_storage_inactive_user_data | integer | query | False | Filter by constituents.space.block_storage_inactive_user_data<br><br>• Introduced in: 9.9 |
| constituents.space.total_metadata | integer | query | False | Filter by constituents.space.total_metadata<br><br>• Introduced in: 9.15 |
| constituents.space.available_percent | integer | query | False | Filter by constituents.space.available_percent<br><br>• Introduced in: 9.9 |
| constituents.space.max_size | string | query | False | Filter by constituents.space.max_size<br><br>• Introduced in: 9.14 |
| constituents.space.afs_total | integer | query | False | Filter by constituents.space.afs_total<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| constituents.space.capacity_tier_footprint | integer | query | False | Filter by constituents.space.capacity_tier_footprint<br><br>• Introduced in: 9.9 |
| constituents.space.local_tier_footprint | integer | query | False | Filter by constituents.space.local_tier_footprint<br><br>• Introduced in: 9.9 |
| constituents.node.uuid | string | query | False | Filter by constituents.node.uuid<br><br>• Introduced in: 9.17 |
| constituents.node.name | string | query | False | Filter by constituents.node.name<br><br>• Introduced in: 9.17 |
| scheduled_snapshot_naming_scheme | string | query | False | Filter by scheduled_snapshot_naming_scheme<br><br>• Introduced in: 9.10 |
| efficiency.last_op_state | string | query | False | Filter by efficiency.last_op_state<br><br>• Introduced in: 9.9 |
| efficiency.last_op_err | string | query | False | Filter by efficiency.last_op_err<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| efficiency.progress | string | query | False | Filter by efficiency.progress<br><br>• Introduced in: 9.9 |
| efficiency.last_op_size | integer | query | False | Filter by efficiency.last_op_size<br><br>• Introduced in: 9.9 |
| efficiency.schedule | string | query | False | Filter by efficiency.schedule<br><br>• Introduced in: 9.8 |
| efficiency.state | string | query | False | Filter by efficiency.state<br><br>• Introduced in: 9.9 |
| efficiency.compaction | string | query | False | Filter by efficiency.compaction |
| efficiency.last_op_begin | string | query | False | Filter by efficiency.last_op_begin<br><br>• Introduced in: 9.9 |
| efficiency.scanner.dedupe | boolean | query | False | Filter by efficiency.scanner.dedupe<br><br>• Introduced in: 9.11 |
| efficiency.scanner.state | string | query | False | Filter by efficiency.scanner.state<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| efficiency.scanner.scan_old_data | boolean | query | False | Filter by efficiency.scanner.scan_old_data<br>• Introduced in: 9.11 |
| efficiency.scanner.compression | boolean | query | False | Filter by efficiency.scanner.compression<br>• Introduced in: 9.11 |
| efficiency.dedupe | string | query | False | Filter by efficiency.dedupe |
| efficiency.logging_enabled | boolean | query | False | Filter by efficiency.logging_enabled<br>• Introduced in: 9.11 |
| efficiency.ratio | number | query | False | Filter by efficiency.ratio<br>• Introduced in: 9.17 |
| efficiency.type | string | query | False | Filter by efficiency.type<br>• Introduced in: 9.9 |
| efficiency.policy.name | string | query | False | Filter by efficiency.policy.name<br>• Introduced in: 9.7 |
| efficiency.auto_state | string | query | False | Filter by efficiency.auto_state<br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| efficiency.compression | string | query | False | Filter by efficiency.compression |
| efficiency.application_io_size | string | query | False | Filter by efficiency.application_io_size<br>• Introduced in: 9.8 |
| efficiency.space_savings.dedupe_sharing | integer | query | False | Filter by efficiency.space_savings.dedupe_sharing<br>• Introduced in: 9.11 |
| efficiency.space_savings.dedupe_percent | integer | query | False | Filter by efficiency.space_savings.dedupe_percent<br>• Introduced in: 9.11 |
| efficiency.space_savings.total_percent | integer | query | False | Filter by efficiency.space_savings.total_percent<br>• Introduced in: 9.11 |
| efficiency.space_savings.compression_percent | integer | query | False | Filter by efficiency.space_savings.compression_percent<br>• Introduced in: 9.11 |
| efficiency.space_savings.compression | integer | query | False | Filter by efficiency.space_savings.compression<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| efficiency.space_savings.total | integer | query | False | Filter by efficiency.space_savings.total<br><br>• Introduced in: 9.11 |
| efficiency.space_savings.dedupe | integer | query | False | Filter by efficiency.space_savings.dedupe<br><br>• Introduced in: 9.11 |
| efficiency.storage_efficiency_mode | string | query | False | Filter by efficiency.storage_efficiency_mode<br><br>• Introduced in: 9.10 |
| efficiency.volume_path | string | query | False | Filter by efficiency.volume_path<br><br>• Introduced in: 9.13 |
| efficiency.has_savings | boolean | query | False | Filter by efficiency.has_savings<br><br>• Introduced in: 9.11 |
| efficiency.last_op_end | string | query | False | Filter by efficiency.last_op_end<br><br>• Introduced in: 9.9 |
| efficiency.compression_type | string | query | False | Filter by efficiency.compression_type<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| efficiency.cross_volume_dedupe | string | query | False | Filter by efficiency.cross_volume_dedupe |
| efficiency.op_state | string | query | False | Filter by efficiency.op_state<br>• Introduced in: 9.9 |
| efficiency.idcs_scanner.operation_state | string | query | False | Filter by efficiency.idcs_scanner.operation_state<br>• Introduced in: 9.13 |
| efficiency.idcs_scanner.status | string | query | False | Filter by efficiency.idcs_scanner.status<br>• Introduced in: 9.13 |
| efficiency.idcs_scanner.threshold_inactive_time | string | query | False | Filter by efficiency.idcs_scanner.threshold_inactive_time<br>• Introduced in: 9.13 |
| efficiency.idcs_scanner.enabled | boolean | query | False | Filter by efficiency.idcs_scanner.enabled<br>• Introduced in: 9.13 |
| efficiency.idcs_scanner.mode | string | query | False | Filter by efficiency.idcs_scanner.mode<br>• Introduced in: 9.13 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| qos.policy.min_throughput_mbps | integer | query | False | Filter by qos.policy.min_throughput_mbps<br>• Introduced in: 9.8<br>• Max value: 4194303<br>• Min value: 0 |
| qos.policy.max_throughput_iops | integer | query | False | Filter by qos.policy.max_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| qos.policy.max_throughput_mbps | integer | query | False | Filter by qos.policy.max_throughput_mbps<br>• Max value: 4194303<br>• Min value: 0 |
| qos.policy.name | string | query | False | Filter by qos.policy.name |
| qos.policy.min_throughput_iops | integer | query | False | Filter by qos.policy.min_throughput_iops<br>• Max value: 2147483647<br>• Min value: 0 |
| qos.policy.uuid | string | query | False | Filter by qos.policy.uuid |
| qos.policy.max_throughput | string | query | False | Filter by qos.policy.max_throughput<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| qos.policy.min_throughput | string | query | False | Filter by qos.policy.min_throughput<br>• Introduced in: 9.17 |
| msid | integer | query | False | Filter by msid<br>• Introduced in: 9.11 |
| cloud_retrieval_policy | string | query | False | Filter by cloud_retrieval_policy<br>• Introduced in: 9.8 |
| _tags | string | query | False | Filter by _tags<br>• Introduced in: 9.13 |
| guarantee.honored | boolean | query | False | Filter by guarantee.honored |
| guarantee.type | string | query | False | Filter by guarantee.type |
| snapshot_locking_enabled | boolean | query | False | Filter by snapshot_locking_enabled<br>• Introduced in: 9.12 |
| activity_tracking.notices.code | string | query | False | Filter by activity_tracking.notices.code<br>• Introduced in: 9.18 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| activity_tracking.notices.message | string | query | False | Filter by activity_tracking.notices.message<br><br>• Introduced in: 9.18 |
| activity_tracking.unsupported_reason.message | string | query | False | Filter by activity_tracking.unsupported_reason.message<br><br>• Introduced in: 9.10 |
| activity_tracking.unsupported_reason.code | string | query | False | Filter by activity_tracking.unsupported_reason.code<br><br>• Introduced in: 9.10 |
| activity_tracking.supported | boolean | query | False | Filter by activity_tracking.supported<br><br>• Introduced in: 9.10 |
| activity_tracking.state | string | query | False | Filter by activity_tracking.state<br><br>• Introduced in: 9.10 |
| space.metadata | integer | query | False | Filter by space.metadata |
| space.available_percent | integer | query | False | Filter by space.available_percent<br><br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.volume_guarantee_footprint | integer | query | False | Filter by space.volume_guarantee_footprint<br>• Introduced in: 9.10 |
| space.filesystem_size | integer | query | False | Filter by space.filesystem_size<br>• Introduced in: 9.10 |
| space.physical_used_percent | integer | query | False | Filter by space.physical_used_percent<br>• Introduced in: 9.10 |
| space.capacity_tier_footprint | integer | query | False | Filter by space.capacity_tier_footprint |
| space.cross_volume_dedupe_metafiles_footprint | integer | query | False | Filter by space.cross_volume_dedupe_metafiles_footprint<br>• Introduced in: 9.10 |
| space.full_threshold_percent | integer | query | False | Filter by space.full_threshold_percent<br>• Introduced in: 9.9 |
| space.overwrite_reserve | integer | query | False | Filter by space.overwrite_reserve<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.logical_space.used | integer | query | False | Filter by space.logical_space.used<br><br>• Introduced in: 9.9 |
| space.logical_space.used_by_snapshots | integer | query | False | Filter by space.logical_space.used_by_snapshots<br><br>• Introduced in: 9.10 |
| space.logical_space.used_percent | integer | query | False | Filter by space.logical_space.used_percent<br><br>• Introduced in: 9.9 |
| space.logical_space.reporting | boolean | query | False | Filter by space.logical_space.reporting |
| space.logical_space.enforcement | boolean | query | False | Filter by space.logical_space.enforcement |
| space.logical_space.available | integer | query | False | Filter by space.logical_space.available |
| space.logical_space.used_by_afs | integer | query | False | Filter by space.logical_space.used_by_afs |
| space.physical_used | integer | query | False | Filter by space.physical_used<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.auto_adaptive_compression_footprint_data_reduction | integer | query | False | Filter by space.auto_adaptive_compression_footprint_data_reduction<br><br>• Introduced in: 9.11 |
| space.block_storage_inactive_user_data_percent | integer | query | False | Filter by space.block_storage_inactive_user_data_percent<br><br>• Introduced in: 9.9 |
| space.used_by_afs | integer | query | False | Filter by space.used_by_afs<br><br>• Introduced in: 9.9 |
| space.is_used_stale | boolean | query | False | Filter by space.is_used_stale<br><br>• Introduced in: 9.12 |
| space.used | integer | query | False | Filter by space.used |
| space.size_available_for_snapshots | integer | query | False | Filter by space.size_available_for_snapshots<br><br>• Introduced in: 9.9 |
| space.cross_volume_dedupe_metafiles_temporary_footprint | integer | query | False | Filter by space.cross_volume_dedupe_metafiles_temporary_footprint<br><br>• Introduced in: 9.10 |
| space.footprint | integer | query | False | Filter by space.footprint |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| space.user_data | integer | query | False | Filter by space.user_data<br>• Introduced in: 9.10 |
| space.over_provisioned | integer | query | False | Filter by space.over_provisioned |
| space.large_size_enabled | boolean | query | False | Filter by space.large_size_enabled<br>• Introduced in: 9.12 |
| space.dedupe_metafiles_temporary_footprint | integer | query | False | Filter by space.dedupe_metafiles_temporary_footprint<br>• Introduced in: 9.10 |
| space.percent_used | integer | query | False | Filter by space.percent_used<br>• Introduced in: 9.9 |
| space.nearly_full_threshold_percent | integer | query | False | Filter by space.nearly_full_threshold_percent<br>• Introduced in: 9.9 |
| space.compaction_footprint_data_reduction | integer | query | False | Filter by space.compaction_footprint_data_reduction<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| space.file_operation_metadata | integer | query | False | Filter by space.file_operation_metadata<br><br>• Introduced in: 9.10 |
| space.snapshot.reserve_size | integer | query | False | Filter by space.snapshot.reserve_size<br><br>• Introduced in: 9.9 |
| space.snapshot.autodelete.commitment | string | query | False | Filter by space.snapshot.autodelete.commitment<br><br>• Introduced in: 9.13 |
| space.snapshot.autodelete.target_free_space | integer | query | False | Filter by space.snapshot.autodelete.target_free_space<br><br>• Introduced in: 9.13 |
| space.snapshot.autodelete.delete_order | string | query | False | Filter by space.snapshot.autodelete.delete_order<br><br>• Introduced in: 9.13 |
| space.snapshot.autodelete.trigger | string | query | False | Filter by space.snapshot.autodelete.trigger<br><br>• Introduced in: 9.13 |
| space.snapshot.autodelete.defer_delete | string | query | False | Filter by space.snapshot.autodelete.defer_delete<br><br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.snapshot.autodelete.prefix | string | query | False | Filter by space.snapshot.autodelete.prefix<br><br>• Introduced in: 9.13 |
| space.snapshot.autodelete_trigger | string | query | False | Filter by space.snapshot.autodelete_trigger<br><br>• Introduced in: 9.10 |
| space.snapshot.reserve_percent | integer | query | False | Filter by space.snapshot.reserve_percent |
| space.snapshot.used | integer | query | False | Filter by space.snapshot.used |
| space.snapshot.space_used_percent | integer | query | False | Filter by space.snapshot.space_used_percent<br><br>• Introduced in: 9.9 |
| space.snapshot.reserve_available | integer | query | False | Filter by space.snapshot.reserve_available<br><br>• Introduced in: 9.10 |
| space.available | integer | query | False | Filter by space.available |
| space.size | integer | query | False | Filter by space.size |
| space.block_storage_inactive_user_data | integer | query | False | Filter by space.block_storage_inactive_user_data |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| space.total_metadata | integer | query | False | Filter by space.total_metadata<br>• Introduced in: 9.15 |
| space.max_size | string | query | False | Filter by space.max_size<br>• Introduced in: 9.14 |
| space.afs_total | integer | query | False | Filter by space.afs_total<br>• Introduced in: 9.9 |
| space.snapmirror_destination_footprint | integer | query | False | Filter by space.snapmirror_destination_footprint<br>• Introduced in: 9.10 |
| space.local_tier_footprint | integer | query | False | Filter by space.local_tier_footprint<br>• Introduced in: 9.8 |
| space.total_metadata_footprint | integer | query | False | Filter by space.total_metadata_footprint<br>• Introduced in: 9.15 |
| space.snapshot_reserve_unusable | integer | query | False | Filter by space.snapshot_reserve_unusable<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.fractional_reserve | integer | query | False | Filter by space.fractional_reserve<br><br>• Introduced in: 9.9 |
| space.dedupe_metafiles_footprint | integer | query | False | Filter by space.dedupe_metafiles_footprint<br><br>• Introduced in: 9.10 |
| space.effective_total_footprint | integer | query | False | Filter by space.effective_total_footprint<br><br>• Introduced in: 9.11 |
| space.performance_tier_footprint | integer | query | False | Filter by space.performance_tier_footprint<br><br>• Introduced in: 9.8 |
| space.filesystem_size_fixed | boolean | query | False | Filter by space.filesystem_size_fixed<br><br>• Introduced in: 9.10 |
| space.total_footprint | integer | query | False | Filter by space.total_footprint<br><br>• Introduced in: 9.8 |
| space.snapshot_spill | integer | query | False | Filter by space.snapshot_spill<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| space.capacity_tier_footprint_data_reduction | integer | query | False | Filter by space.capacity_tier_footprint_data_reduction<br>• Introduced in: 9.13 |
| space.overwrite_reserve_used | integer | query | False | Filter by space.overwrite_reserve_used<br>• Introduced in: 9.9 |
| space.delayed_free_footprint | integer | query | False | Filter by space.delayed_free_footprint<br>• Introduced in: 9.10 |
| space.expected_available | integer | query | False | Filter by space.expected_available<br>• Introduced in: 9.10 |
| uuid | string | query | False | Filter by uuid |
| tiering.min_cooling_days | integer | query | False | Filter by tiering.min_cooling_days<br>• Introduced in: 9.8<br>• Max value: 183<br>• Min value: 2 |
| tiering.policy | string | query | False | Filter by tiering.policy |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| tiering.object_tags | string | query | False | Filter by tiering.object_tags<br>• Introduced in: 9.8<br>• maxLength: 257 |
| convert_unicode | boolean | query | False | Filter by convert_unicode<br>• Introduced in: 9.10 |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other |
| statistics.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.flexcache_raw.client_requested_blocks | integer | query | False | Filter by statistics.flexcache_raw.client_requested_blocks<br><br>• Introduced in: 9.8 |
| statistics.flexcache_raw.cache_miss_blocks | integer | query | False | Filter by statistics.flexcache_raw.cache_miss_blocks<br><br>• Introduced in: 9.8 |
| statistics.flexcache_raw.status | string | query | False | Filter by statistics.flexcache_raw.status<br><br>• Introduced in: 9.8 |
| statistics.flexcache_raw.timestamp | string | query | False | Filter by statistics.flexcache_raw.timestamp<br><br>• Introduced in: 9.8 |
| statistics.cloud.status | string | query | False | Filter by statistics.cloud.status<br><br>• Introduced in: 9.7 |
| statistics.cloud.latency_raw.other | integer | query | False | Filter by statistics.cloud.latency_raw.other<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cloud.latency_raw.total | integer | query | False | Filter by statistics.cloud.latency_raw.total<br><br>• Introduced in: 9.7 |
| statistics.cloud.latency_raw.write | integer | query | False | Filter by statistics.cloud.latency_raw.write<br><br>• Introduced in: 9.7 |
| statistics.cloud.latency_raw.read | integer | query | False | Filter by statistics.cloud.latency_raw.read<br><br>• Introduced in: 9.7 |
| statistics.cloud.iops_raw.other | integer | query | False | Filter by statistics.cloud.iops_raw.other<br><br>• Introduced in: 9.7 |
| statistics.cloud.iops_raw.total | integer | query | False | Filter by statistics.cloud.iops_raw.total<br><br>• Introduced in: 9.7 |
| statistics.cloud.iops_raw.write | integer | query | False | Filter by statistics.cloud.iops_raw.write<br><br>• Introduced in: 9.7 |
| statistics.cloud.iops_raw.read | integer | query | False | Filter by statistics.cloud.iops_raw.read<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.cloud.timestamp | string | query | False | Filter by statistics.cloud.timestamp<br><br>• Introduced in: 9.7 |
| statistics.nfs_ops_raw.unlink.count | integer | query | False | Filter by statistics.nfs_ops_raw.unlink.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.unlink.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.unlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.lookup.count | integer | query | False | Filter by statistics.nfs_ops_raw.lookup.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.lookup.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.lookup.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.open.count | integer | query | False | Filter by statistics.nfs_ops_raw.open.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.open.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.open.total_time<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.nfs_ops_raw.access.count | integer | query | False | Filter by statistics.nfs_ops_raw.access.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.access.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.access.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.link.count | integer | query | False | Filter by statistics.nfs_ops_raw.link.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.link.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.link.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.lock.count | integer | query | False | Filter by statistics.nfs_ops_raw.lock.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.lock.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.lock.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.rename.count | integer | query | False | Filter by statistics.nfs_ops_raw.rename.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.nfs_ops_raw.rename.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.rename.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.getattr.count | integer | query | False | Filter by statistics.nfs_ops_raw.getattr.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.getattr.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.getattr.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels | string | query | False | Filter by statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.read.count | integer | query | False | Filter by statistics.nfs_ops_raw.read.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_counts | integer | query | False | Filter by statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_counts<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.nfs_ops_raw.read.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.read.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.read.volume_protocol_size_histogram_labels | string | query | False | Filter by statistics.nfs_ops_raw.read.volume_protocol_size_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts | integer | query | False | Filter by statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels | string | query | False | Filter by statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.write.count | integer | query | False | Filter by statistics.nfs_ops_raw.write.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_counts | integer | query | False | Filter by statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_counts<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.nfs_ops_raw.write.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.write.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.write.volume_protocol_size_histogram_labels | string | query | False | Filter by statistics.nfs_ops_raw.write.volume_protocol_size_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts | integer | query | False | Filter by statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.readdir.count | integer | query | False | Filter by statistics.nfs_ops_raw.readdir.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.readdir.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.readdir.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.setattr.count | integer | query | False | Filter by statistics.nfs_ops_raw.setattr.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.nfs_ops_raw.setattr.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.setattr.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.watch.count | integer | query | False | Filter by statistics.nfs_ops_raw.watch.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.watch.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.watch.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.file.count | integer | query | False | Filter by statistics.nfs_ops_raw.create.file.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.file.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.create.file.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.symlink.count | integer | query | False | Filter by statistics.nfs_ops_raw.create.symlink.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.nfs_ops_raw.create.symlink.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.create.symlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.dir.count | integer | query | False | Filter by statistics.nfs_ops_raw.create.dir.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.dir.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.create.dir.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.other.count | integer | query | False | Filter by statistics.nfs_ops_raw.create.other.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.create.other.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.create.other.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.audit.count | integer | query | False | Filter by statistics.nfs_ops_raw.audit.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.nfs_ops_raw.audit.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.audit.total_time<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.readlink.count | integer | query | False | Filter by statistics.nfs_ops_raw.readlink.count<br><br>• Introduced in: 9.11 |
| statistics.nfs_ops_raw.readlink.total_time | integer | query | False | Filter by statistics.nfs_ops_raw.readlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.throughput_raw.other | integer | query | False | Filter by statistics.throughput_raw.other |
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write |
| statistics.throughput_raw.read | integer | query | False | Filter by statistics.throughput_raw.read |
| statistics.status | string | query | False | Filter by statistics.status |
| statistics.cifs_ops_raw.unlink.count | integer | query | False | Filter by statistics.cifs_ops_raw.unlink.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.unlink.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.unlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.lookup.count | integer | query | False | Filter by statistics.cifs_ops_raw.lookup.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.lookup.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.lookup.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.open.count | integer | query | False | Filter by statistics.cifs_ops_raw.open.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.open.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.open.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.access.count | integer | query | False | Filter by statistics.cifs_ops_raw.access.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.access.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.access.total_time<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.link.count | integer | query | False | Filter by statistics.cifs_ops_raw.link.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.link.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.link.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.lock.count | integer | query | False | Filter by statistics.cifs_ops_raw.lock.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.lock.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.lock.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.rename.count | integer | query | False | Filter by statistics.cifs_ops_raw.rename.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.rename.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.rename.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.getattr.count | integer | query | False | Filter by statistics.cifs_ops_raw.getattr.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.getattr.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.getattr.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels | string | query | False | Filter by statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.read.count | integer | query | False | Filter by statistics.cifs_ops_raw.read.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_counts | integer | query | False | Filter by statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.read.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.read.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.read.volume_protocol_size_histogram_labels | string | query | False | Filter by statistics.cifs_ops_raw.read.volume_protocol_size_histogram_labels<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts | integer | query | False | Filter by statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels | string | query | False | Filter by statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.write.count | integer | query | False | Filter by statistics.cifs_ops_raw.write.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_counts | integer | query | False | Filter by statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.write.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.write.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.write.volume_protocol_size_histogram_labels | string | query | False | Filter by statistics.cifs_ops_raw.write.volume_protocol_size_histogram_labels<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts | integer | query | False | Filter by statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.readdir.count | integer | query | False | Filter by statistics.cifs_ops_raw.readdir.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.readdir.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.readdir.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.setattr.count | integer | query | False | Filter by statistics.cifs_ops_raw.setattr.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.setattr.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.setattr.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.watch.count | integer | query | False | Filter by statistics.cifs_ops_raw.watch.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.watch.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.watch.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.file.count | integer | query | False | Filter by statistics.cifs_ops_raw.create.file.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.file.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.create.file.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.symlink.count | integer | query | False | Filter by statistics.cifs_ops_raw.create.symlink.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.symlink.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.create.symlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.dir.count | integer | query | False | Filter by statistics.cifs_ops_raw.create.dir.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| statistics.cifs_ops_raw.create.dir.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.create.dir.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.other.count | integer | query | False | Filter by statistics.cifs_ops_raw.create.other.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.create.other.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.create.other.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.audit.count | integer | query | False | Filter by statistics.cifs_ops_raw.audit.count<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.audit.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.audit.total_time<br><br>• Introduced in: 9.11 |
| statistics.cifs_ops_raw.readlink.count | integer | query | False | Filter by statistics.cifs_ops_raw.readlink.count<br><br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| statistics.cifs_ops_raw.readlink.total_time | integer | query | False | Filter by statistics.cifs_ops_raw.readlink.total_time<br><br>• Introduced in: 9.11 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp |
| type | string | query | False | Filter by type |
| validate_only | boolean | query | False | Filter by validate_only<br><br>• Introduced in: 9.14 |
| flash_pool.cache_eligibility | string | query | False | Filter by flash_pool.cache_eligibility<br><br>• Introduced in: 9.10 |
| flash_pool.caching_policy | string | query | False | Filter by flash_pool.caching_policy<br><br>• Introduced in: 9.10 |
| flash_pool.cache_retention_priority | string | query | False | Filter by flash_pool.cache_retention_priority<br><br>• Introduced in: 9.10 |
| max_dir_size | integer | query | False | Filter by max_dir_size<br><br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cloud_write_enabled | boolean | query | False | Filter by cloud_write_enabled<br>• Introduced in: 9.13 |
| movement.tiering_policy | string | query | False | Filter by movement.tiering_policy<br>• Introduced in: 9.18 |
| movement.percent_complete | integer | query | False | Filter by movement.percent_complete |
| movement.state | string | query | False | Filter by movement.state |
| movement.cutover_window | integer | query | False | Filter by movement.cutover_window |
| movement.start_time | string | query | False | Filter by movement.start_time<br>• Introduced in: 9.9 |
| movement.capacity_tier_optimized | boolean | query | False | Filter by movement.capacity_tier_optimized<br>• Introduced in: 9.18 |
| movement.destination_aggregate.name | string | query | False | Filter by movement.destination_aggregate.name |
| movement.destination_aggregate.uuid | string | query | False | Filter by movement.destination_aggregate.uuid |
| encryption.key_id | string | query | False | Filter by encryption.key_id |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| encryption.key_create_time | string | query | False | Filter by encryption.key_create_time<br>• Introduced in: 9.11 |
| encryption.status.message | string | query | False | Filter by encryption.status.message |
| encryption.status.code | string | query | False | Filter by encryption.status.code |
| encryption.action | string | query | False | Filter by encryption.action<br>• Introduced in: 9.14 |
| encryption.enabled | boolean | query | False | Filter by encryption.enabled |
| encryption.type | string | query | False | Filter by encryption.type |
| encryption.rekey | boolean | query | False | Filter by encryption.rekey |
| encryption.state | string | query | False | Filter by encryption.state |
| status | string | query | False | Filter by status<br>• Introduced in: 9.9 |
| snaplock.type | string | query | False | Filter by snaplock.type |
| snaplock.is_audit_log | boolean | query | False | Filter by snaplock.is_audit_log |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| snaplock.privileged_delete | string | query | False | Filter by snaplock.privileged_delete |
| snaplock.retention.minimum | string | query | False | Filter by snaplock.retention.minimum |
| snaplock.retention.default | string | query | False | Filter by snaplock.retention.default |
| snaplock.retention.maximum | string | query | False | Filter by snaplock.retention.maximum |
| snaplock.compliance_clock_time | string | query | False | Filter by snaplock.compliance_clock_time |
| snaplock.expiry_time | string | query | False | Filter by snaplock.expiry_time |
| snaplock.append_mode_enabled | boolean | query | False | Filter by snaplock.append_mode_enabled |
| snaplock.autocommit_period | string | query | False | Filter by snaplock.autocommit_period |
| snaplock.unspecified_retention_file_count | integer | query | False | Filter by snaplock.unspecified_retention_file_count<br>• Introduced in: 9.8 |
| snaplock.litigation_count | integer | query | False | Filter by snaplock.litigation_count |
| nas.gid | integer | query | False | Filter by nas.gid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| nas.export_policy.name | string | query | False | Filter by nas.export_policy.name |
| nas.export_policy.id | integer | query | False | Filter by nas.export_policy.id |
| nas.unix_permissions | integer | query | False | Filter by nas.unix_permissions |
| nas.path | string | query | False | Filter by nas.path |
| nas.junction_parent.name | string | query | False | Filter by nas.junction_parent.name<br>• Introduced in: 9.9 |
| nas.junction_parent.uuid | string | query | False | Filter by nas.junction_parent.uuid<br>• Introduced in: 9.9 |
| nas.security_style | string | query | False | Filter by nas.security_style |
| nas.uid | integer | query | False | Filter by nas.uid |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total |
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.timestamp | string | query | False | Filter by metric.timestamp |
| metric.iops.other | integer | query | False | Filter by metric.iops.other |
| metric.iops.total | integer | query | False | Filter by metric.iops.total |
| metric.iops.write | integer | query | False | Filter by metric.iops.write |
| metric.iops.read | integer | query | False | Filter by metric.iops.read |
| metric.latency.other | integer | query | False | Filter by metric.latency.other |
| metric.latency.total | integer | query | False | Filter by metric.latency.total |
| metric.latency.write | integer | query | False | Filter by metric.latency.write |
| metric.latency.read | integer | query | False | Filter by metric.latency.read |
| metric.status | string | query | False | Filter by metric.status |
| metric.cloud.duration | string | query | False | Filter by metric.cloud.duration<br>• Introduced in: 9.7 |
| metric.cloud.status | string | query | False | Filter by metric.cloud.status<br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.cloud.timestamp | string | query | False | Filter by metric.cloud.timestamp<br><br>• Introduced in: 9.7 |
| metric.cloud.iops.other | integer | query | False | Filter by metric.cloud.iops.other<br><br>• Introduced in: 9.7 |
| metric.cloud.iops.total | integer | query | False | Filter by metric.cloud.iops.total<br><br>• Introduced in: 9.7 |
| metric.cloud.iops.write | integer | query | False | Filter by metric.cloud.iops.write<br><br>• Introduced in: 9.7 |
| metric.cloud.iops.read | integer | query | False | Filter by metric.cloud.iops.read<br><br>• Introduced in: 9.7 |
| metric.cloud.latency.other | integer | query | False | Filter by metric.cloud.latency.other<br><br>• Introduced in: 9.7 |
| metric.cloud.latency.total | integer | query | False | Filter by metric.cloud.latency.total<br><br>• Introduced in: 9.7 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| metric.cloud.latency.write | integer | query | False | Filter by metric.cloud.latency.write<br><br>• Introduced in: 9.7 |
| metric.cloud.latency.read | integer | query | False | Filter by metric.cloud.latency.read<br><br>• Introduced in: 9.7 |
| metric.flexcache.timestamp | string | query | False | Filter by metric.flexcache.timestamp<br><br>• Introduced in: 9.8 |
| metric.flexcache.cache_miss_percent | integer | query | False | Filter by metric.flexcache.cache_miss_percent<br><br>• Introduced in: 9.8 |
| metric.flexcache.bandwidth_savings | integer | query | False | Filter by metric.flexcache.bandwidth_savings<br><br>• Introduced in: 9.9 |
| metric.flexcache.duration | string | query | False | Filter by metric.flexcache.duration<br><br>• Introduced in: 9.8 |
| metric.flexcache.status | string | query | False | Filter by metric.flexcache.status<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| metric.duration | string | query | False | Filter by metric.duration |
| flexcache_endpoint_type | string | query | False | Filter by flexcache_endpoint_type |
| analytics.unsupported_reason.code | string | query | False | Filter by analytics.unsupported_reason.code<br>• Introduced in: 9.8 |
| analytics.unsupported_reason.message | string | query | False | Filter by analytics.unsupported_reason.message<br>• Introduced in: 9.8 |
| analytics.total_files | integer | query | False | Filter by analytics.total_files<br>• Introduced in: 9.14 |
| analytics.by_modified_time.bytes_used.labels | string | query | False | Filter by analytics.by_modified_time.bytes_used.labels<br>• Introduced in: 9.17 |
| analytics.by_modified_time.bytes_used.newest_label | string | query | False | Filter by analytics.by_modified_time.bytes_used.newest_label<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| analytics.by_modified_time.bytes_used.oldest_label | string | query | False | Filter by analytics.by_modified_time.bytes_used.oldest_label<br><br>• Introduced in: 9.17 |
| analytics.by_modified_time.bytes_used.values | integer | query | False | Filter by analytics.by_modified_time.bytes_used.values<br><br>• Introduced in: 9.17 |
| analytics.by_modified_time.bytes_used.aged_data_metric | number | query | False | Filter by analytics.by_modified_time.bytes_used.aged_data_metric<br><br>• Introduced in: 9.17 |
| analytics.by_modified_time.bytes_used.percentages | number | query | False | Filter by analytics.by_modified_time.bytes_used.percentages<br><br>• Introduced in: 9.17 |
| analytics.scan_progress | integer | query | False | Filter by analytics.scan_progress<br><br>• Introduced in: 9.8 |
| analytics.initialization.state | string | query | False | Filter by analytics.initialization.state<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| analytics.bytes_used | integer | query | False | Filter by analytics.bytes_used<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.labels | string | query | False | Filter by analytics.by_accessed_time.bytes_used.labels<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.newest_label | string | query | False | Filter by analytics.by_accessed_time.bytes_used.newest_label<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.oldest_label | string | query | False | Filter by analytics.by_accessed_time.bytes_used.oldest_label<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.values | integer | query | False | Filter by analytics.by_accessed_time.bytes_used.values<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.aged_data_metric | number | query | False | Filter by analytics.by_accessed_time.bytes_used.aged_data_metric<br><br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| analytics.by_accessed_time.bytes_used.percentages | number | query | False | Filter by analytics.by_accessed_time.bytes_used.percentages<br><br>• Introduced in: 9.17 |
| analytics.report_time | string | query | False | Filter by analytics.report_time<br><br>• Introduced in: 9.17 |
| analytics.scan_throttle_reason.message | string | query | False | Filter by analytics.scan_throttle_reason.message<br><br>• Introduced in: 9.14 |
| analytics.scan_throttle_reason.arguments | string | query | False | Filter by analytics.scan_throttle_reason.arguments<br><br>• Introduced in: 9.14 |
| analytics.scan_throttle_reason.code | string | query | False | Filter by analytics.scan_throttle_reason.code<br><br>• Introduced in: 9.14 |
| analytics.subdir_count | integer | query | False | Filter by analytics.subdir_count<br><br>• Introduced in: 9.17 |
| analytics.files_scanned | integer | query | False | Filter by analytics.files_scanned<br><br>• Introduced in: 9.14 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| analytics.file_count | integer | query | False | Filter by analytics.file_count<br><br>• Introduced in: 9.17 |
| analytics.state | string | query | False | Filter by analytics.state<br><br>• Introduced in: 9.8 |
| analytics.supported | boolean | query | False | Filter by analytics.supported<br><br>• Introduced in: 9.8 |
| analytics.incomplete_data | boolean | query | False | Filter by analytics.incomplete_data<br><br>• Introduced in: 9.17 |
| snapshot_directory_access_enabled | boolean | query | False | Filter by snapshot_directory_access_enabled<br><br>• Introduced in: 9.13 |
| state | string | query | False | Filter by state |
| is_large_dir_enabled | boolean | query | False | Filter by is_large_dir_enabled<br><br>• Introduced in: 9.18 |
| aggregates.name | string | query | False | Filter by aggregates.name |
| aggregates.uuid | string | query | False | Filter by aggregates.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| autosize.shrink_threshold | integer | query | False | Filter by autosize.shrink_threshold |
| autosize.grow_threshold | integer | query | False | Filter by autosize.grow_threshold |
| autosize.mode | string | query | False | Filter by autosize.mode |
| autosize.maximum | integer | query | False | Filter by autosize.maximum |
| autosize.minimum | integer | query | False | Filter by autosize.minimum |
| snapshot_count | integer | query | False | Filter by snapshot_count<br>• Introduced in: 9.10<br>• Max value: 1023<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| analytics | analytics | Additional file system analytics information that is non-modifiable amongst all elements in the collection.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>This analytics object captures properties that are non-modifiable amongst all elements included in the `records` array. The non-modifiable properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large. |
| num_records | integer | Number of records |
| records | array[volume] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
      },
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "activity_tracking": {
        "notices": [
          {
            "code": "124518424",
            "message": "The top metrics report contains partial data
for read operations because NFSv4 reads using Multi-Processor I/O
(MPIO) are not tracked."
          }
        ],
        "state": "string",
        "unsupported_reason": {
          "code": "124518405",
          "message": "Volume activity tracking cannot be enabled on
volumes that contain LUNs."
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "aggressive_readahead_mode": "string",
      "analytics": {
        "by_accessed_time": {
          "bytes_used": {
            "aged_data_metric": 15.23,
            "labels": [
              "2019-07",
              "2019-06",
              "2019-05",
              "2019",
              "2018",
              "--2017",
              "unknown"
            ],
```

```
            "newest_label": "2019-07",
            "oldest_label": "2019-07",
            "percentages": [
              0.1,
              11.24,
              0.18,
              15.75,
              0.75,
              83.5,
              0
            ],
            "values": [
              15925248,
              1735569408,
              27672576,
              2430595072,
              116105216,
              12889948160,
              0
            ]
          }
        },
        "by_modified_time": {
          "bytes_used": {
            "aged_data_metric": 15.23,
            "labels": [
              "2019-07",
              "2019-06",
              "2019-05",
              "2019",
              "2018",
              "--2017",
              "unknown"
            ],
            "newest_label": "2019-07",
            "oldest_label": "2019-07",
            "percentages": [
              0.1,
              11.24,
              0.18,
              15.75,
              0.75,
              83.5,
              0
            ],
            "values": [
```

```
              15925248,
              1735569408,
              27672576,
              2430595072,
              116105216,
              12889948160,
              0
            ]
          }
        },
        "bytes_used": 15436648448,
        "file_count": 21134,
        "files_scanned": 43002,
        "initialization": {
          "state": "string"
        },
        "report_time": "2024-11-06 13:57:15 -0500",
        "scan_progress": 17,
        "scan_throttle_reason": {
          "arguments": [
            "string"
          ],
          "code": "6739881",
          "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
        },
        "state": "string",
        "subdir_count": 35,
        "total_files": 101890,
        "unsupported_reason": {
          "code": "111411207",
          "message": "File system analytics cannot be enabled on
volumes that contain LUNs."
        }
      },
      "anti_ransomware": {
        "attack_detected_by": "string",
        "attack_detection_parameters": {
          "file_create_op_rate_surge_notify_percent": 100
        },
        "attack_probability": "string",
        "attack_reports": [
          {
            "_links": {
              "suspects": {
```

```
                "href": "/api/resourcelink"
              }
            },
            "time": "2021-06-01 20:36:41 +0530"
          }
        ],
        "block_device_detection_start_time": "string",
        "block_device_detection_state": "string",
        "clear_suspect": {
          "phase": "string",
          "start_time": "string"
        },
        "dry_run_start_time": "string",
        "space": {
          "snapshot_count": 0,
          "used": 0,
          "used_by_logs": 0,
          "used_by_snapshots": 0
        },
        "state": "string",
        "surge_usage": {
          "file_create_peak_rate_per_minute": 10,
          "file_delete_peak_rate_per_minute": 50,
          "file_rename_peak_rate_per_minute": 30,
          "high_entropy_data_write_peak_percent": 30,
          "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
          "time": "2021-12-01 23:16:20 +0530"
        },
        "suspect_files": [
          {
            "count": 0,
            "entropy": "string",
            "format": "string"
          }
        ],
        "typical_usage": {
          "file_create_peak_rate_per_minute": 50,
          "file_delete_peak_rate_per_minute": 10,
          "file_rename_peak_rate_per_minute": 5,
          "high_entropy_data_write_peak_percent": 10,
          "high_entropy_data_write_peak_rate_kb_per_minute": 1200
        },
        "workload": {
          "file_extension_types_count": 3,
          "file_extensions_observed": [
            "pdf",
```

```
            "jpeg",
            "txt"
          ],
          "historical_statistics": {
            "file_create_peak_rate_per_minute": 50,
            "file_delete_peak_rate_per_minute": 10,
            "file_rename_peak_rate_per_minute": 5,
            "high_entropy_data_write_peak_percent": 10,
            "high_entropy_data_write_peak_rate_kb_per_minute": 1200
          },
          "newly_observed_file_extensions": [
            {
              "count": [
                20
              ],
              "name": [
                "lockile"
              ]
            }
          ],
          "surge_statistics": {
            "file_create_peak_rate_per_minute": 10,
            "file_delete_peak_rate_per_minute": 50,
            "file_rename_peak_rate_per_minute": 30,
            "high_entropy_data_write_peak_percent": 30,
            "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
            "time": "2021-12-01 23:16:20 +0530"
          },
          "surge_usage": {
            "file_create_peak_rate_per_minute": 10,
            "file_delete_peak_rate_per_minute": 50,
            "file_rename_peak_rate_per_minute": 30,
            "high_entropy_data_write_peak_percent": 30,
            "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
            "newly_observed_file_extensions": [
              {
                "count": [
                  20
                ],
                "name": [
                  "lockile"
                ]
              }
            ],
            "time": "2021-12-01 23:16:20 +0530"
          },
```

```
          "typical_usage": {
            "file_create_peak_rate_per_minute": 50,
            "file_delete_peak_rate_per_minute": 10,
            "file_rename_peak_rate_per_minute": 5,
            "high_entropy_data_write_peak_percent": 10,
            "high_entropy_data_write_peak_rate_kb_per_minute": 1200
          }
        }
      },
      "application": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
      },
      "asynchronous_directory_delete": {
        "trash_bin": "string"
      },
      "autosize": {
        "mode": "string"
      },
      "clone": {
        "inherited_physical_used": 0,
        "inherited_savings": 0,
        "lun_name": "string",
        "parent_snapshot": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "this_snapshot",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "parent_svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "parent_volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
```

```
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        },
        "qtree_name": "string",
        "split_complete_percent": 0,
        "split_estimate": 0
      },
      "cloud_retrieval_policy": "string",
      "comment": "string",
      "consistency_group": {
        "name": "consistency_group_1",
        "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
      },
      "constituents": [
        {
          "aggregates": {
            "name": "string",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          },
          "movement": {
            "cutover_window": 30,
            "destination_aggregate": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "aggr1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "percent_complete": 0,
            "state": "replicating",
            "tiering_policy": "string"
          },
          "name": "string",
          "node": {
            "name": "string",
            "uuid": "string"
          },
          "space": {
            "available": 0,
            "block_storage_inactive_user_data": 0,
            "capacity_tier_footprint": 0,
            "footprint": 0,
            "local_tier_footprint": 0,
```

```
            "logical_space": {
              "available": 0,
              "used_by_afs": 0
            },
            "max_size": "string",
            "metadata": 0,
            "over_provisioned": 0,
            "performance_tier_footprint": 0,
            "snapshot": {
              "used": 0
            },
            "total_footprint": 0,
            "total_metadata": 0,
            "total_metadata_footprint": 0,
            "used": 0
          }
        }
      ],
      "create_time": "2018-06-04 19:00:00 +0000",
      "efficiency": {
        "application_io_size": "string",
        "auto_state": "string",
        "compaction": "string",
        "compression": "string",
        "compression_type": "string",
        "cross_volume_dedupe": "string",
        "dedupe": "string",
        "idcs_scanner": {
          "mode": "string",
          "operation_state": "string",
          "status": "string",
          "threshold_inactive_time": "P14D"
        },
        "last_op_begin": "string",
        "last_op_end": "string",
        "last_op_err": "string",
        "last_op_size": 0,
        "last_op_state": "string",
        "op_state": "string",
        "policy": {
          "name": "string"
        },
        "progress": "string",
        "ratio": 0,
        "scanner": {
          "state": "string"
```

```
        },
        "schedule": "string",
        "space_savings": {
          "compression": 0,
          "compression_percent": 0,
          "dedupe": 0,
          "dedupe_percent": 0,
          "dedupe_sharing": 0,
          "total": 0,
          "total_percent": 0
        },
        "state": "string",
        "storage_efficiency_mode": "string",
        "type": "string",
        "volume_path": "string"
      },
      "encryption": {
        "action": "string",
        "key_create_time": "2022-01-01 19:00:00 +0000",
        "key_id": "string",
        "state": "string",
        "status": {
          "code": "string",
          "message": "string"
        },
        "type": "string"
      },
      "files": {
        "inodefile_capacity": 0,
        "used": 0
      },
      "flash_pool": {
        "cache_eligibility": "string",
        "cache_retention_priority": "string",
        "caching_policy": "string"
      },
      "flexcache_endpoint_type": "string",
      "flexgroup": {
        "name": "my_flexgroup",
        "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
      },
      "granular_data_mode": "string",
      "guarantee": {
        "type": "string"
      },
      "language": "string",
```

```
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "cloud": {
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "duration": "PT15S",
        "flexcache": {
          "bandwidth_savings": 4096,
          "cache_miss_percent": 20,
          "duration": "PT1D",
          "status": "ok",
          "timestamp": "2017-01-25 11:20:13 +0000"
        },
        "iops": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 11:20:13 +0000"
```

```
      },
      "movement": {
        "cutover_window": 30,
        "destination_aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "start_time": "2020-12-07 03:45:12 -0500",
        "state": "replicating",
        "tiering_policy": "string"
      },
      "name": "vol_cs_dept",
      "nas": {
        "export_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "id": 100,
          "name": "default"
        },
        "junction_parent": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "vs1_root",
          "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
        },
        "path": "/user/my_volume",
        "security_style": "string",
        "unix_permissions": 493
      },
      "nodes": [
        {
          "name": "string",
          "uuid": "string"
        }
```

```
      ],
      "optimize_aggregates": true,
      "qos": {
        "policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "max_throughput": [
            "900KB/s",
            "500MB/s",
            "120GB/s",
            "5000IOPS",
            "5000IOPS,500KB/s",
            "2500IOPS,100MB/s",
            "1000IOPS,25MB/s"
          ],
          "max_throughput_iops": 10000,
          "max_throughput_mbps": 500,
          "min_throughput": [
            "900KB/s",
            "500MB/s",
            "120GB/s",
            "5000IOPS",
            "5000IOPS,500KB/s",
            "2500IOPS,100MB/s",
            "1000IOPS,25MB/s"
          ],
          "min_throughput_iops": 2000,
          "min_throughput_mbps": 500,
          "name": "performance",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "quota": {
        "state": "string"
      },
      "rebalancing": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "data_moved": 0,
        "engine": {
```

```
          "movement": {
            "file_moves_started": 0,
            "last_error": {
              "code": 0,
              "destination": 0,
              "file_id": 0,
              "time": "2018-06-04 19:00:00 +0000"
            },
            "most_recent_start_time": "2018-06-04 19:00:00 +0000"
          },
          "scanner": {
            "blocks_scanned": 0,
            "blocks_skipped": {
              "efficiency_blocks": 0,
              "efficiency_percent": 0,
              "fast_truncate": 0,
              "footprint_invalid": 0,
              "in_snapshot": 0,
              "incompatible": 0,
              "metadata": 0,
              "on_demand_destination": 0,
              "other": 0,
              "remote_cache": 0,
              "too_large": 0,
              "too_small": 0,
              "write_fenced": 0
            },
            "files_scanned": 0,
            "files_skipped": {
              "efficiency_blocks": 0,
              "efficiency_percent": 0,
              "fast_truncate": 0,
              "footprint_invalid": 0,
              "in_snapshot": 0,
              "incompatible": 0,
              "metadata": 0,
              "on_demand_destination": 0,
              "other": 0,
              "remote_cache": 0,
              "too_large": 0,
              "too_small": 0,
              "write_fenced": 0
            }
          }
        },
        "imbalance_percent": 0,
```

```
        "imbalance_size": 0,
        "max_constituent_imbalance_percent": 0,
        "max_runtime": "string",
        "notices": [
          {
            "arguments": [
              {
                "code": "string",
                "message": "string"
              }
            ],
            "code": "4",
            "message": "entry doesn't exist"
          }
        ],
        "runtime": "string",
        "start_time": "string",
        "state": "rebalancing",
        "stop_time": "string",
        "target_used": 0,
        "used_for_imbalance": 0
      },
      "scheduled_snapshot_naming_scheme": "string",
      "sizing_method": "string",
      "snaplock": {
        "append_mode_enabled": "",
        "autocommit_period": "P30M",
        "compliance_clock_time": "2018-06-04 19:00:00 +0000",
        "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
        "is_audit_log": 1,
        "litigation_count": 10,
        "privileged_delete": "enabled",
        "retention": {
          "default": "P30Y",
          "maximum": "P30Y",
          "minimum": "P30Y"
        },
        "type": "enterprise",
        "unspecified_retention_file_count": 10
      },
      "snapshot_count": 0,
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
```

```
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "space": {
        "auto_adaptive_compression_footprint_data_reduction": 0,
        "available": 0,
        "block_storage_inactive_user_data": 0,
        "block_storage_inactive_user_data_percent": 0,
        "capacity_tier_footprint": 0,
        "capacity_tier_footprint_data_reduction": 0,
        "compaction_footprint_data_reduction": 0,
        "cross_volume_dedupe_metafiles_footprint": 0,
        "cross_volume_dedupe_metafiles_temporary_footprint": 0,
        "dedupe_metafiles_footprint": 0,
        "dedupe_metafiles_temporary_footprint": 0,
        "delayed_free_footprint": 0,
        "effective_total_footprint": 0,
        "file_operation_metadata": 0,
        "filesystem_size": 0,
        "footprint": 0,
        "local_tier_footprint": 0,
        "logical_space": {
          "available": 0,
          "used": 0,
          "used_by_afs": 0,
          "used_by_snapshots": 0,
          "used_percent": 0
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "overwrite_reserve": 0,
        "overwrite_reserve_used": 0,
        "percent_used": 0,
        "performance_tier_footprint": 0,
        "size_available_for_snapshots": 0,
        "snapmirror_destination_footprint": 0,
        "snapshot": {
          "autodelete": {
            "commitment": "string",
            "defer_delete": "string",
            "delete_order": "string",
            "prefix": "string",
            "trigger": "string"
          },
```

```
          "autodelete_trigger": "string",
          "reserve_available": 0,
          "reserve_size": 0,
          "space_used_percent": 0,
          "used": 0
        },
        "snapshot_reserve_unusable": 0,
        "snapshot_spill": 0,
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0,
        "user_data": 0,
        "volume_guarantee_footprint": 0
      },
      "state": "string",
      "statistics": {
        "cifs_ops_raw": {
          "access": {
            "count": 1000,
            "total_time": 200
          },
          "audit": {
            "count": 1000,
            "total_time": 200
          },
          "create": {
            "dir": {
              "count": 1000,
              "total_time": 200
            },
            "file": {
              "count": 1000,
              "total_time": 200
            },
            "other": {
              "count": 1000,
              "total_time": 200
            },
            "symlink": {
              "count": 1000,
              "total_time": 200
            }
          },
          "getattr": {
            "count": 1000,
```

```
            "total_time": 200
          },
          "link": {
            "count": 1000,
            "total_time": 200
          },
          "lock": {
            "count": 1000,
            "total_time": 200
          },
          "lookup": {
            "count": 1000,
            "total_time": 200
          },
          "open": {
            "count": 1000,
            "total_time": 200
          },
          "read": {
            "count": 1000,
            "total_time": 200,
            "volume_protocol_latency_histogram_counts": [
              0,
              0,
              0,
              0,
              0,
              15,
              35,
              100,
              200,
              200,
              300,
              500,
              500,
              500,
              1000,
              1000,
              800,
              500,
              500,
              300,
              200,
              50,
              40,
              15,
```

```
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
```

```
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
```

```
            ]
          },
          "readdir": {
            "count": 1000,
            "total_time": 200
          },
          "readlink": {
            "count": 1000,
            "total_time": 200
          },
          "rename": {
            "count": 1000,
            "total_time": 200
          },
          "setattr": {
            "count": 1000,
            "total_time": 200
          },
          "unlink": {
            "count": 1000,
            "total_time": 200
          },
          "watch": {
            "count": 1000,
            "total_time": 200
          },
          "write": {
            "count": 1000,
            "total_time": 200,
            "volume_protocol_latency_histogram_counts": [
              0,
              0,
              0,
              0,
              0,
              15,
              35,
              100,
              200,
              200,
              300,
              500,
              500,
              500,
              1000,
              1000,
```

```
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
```

```
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
```

```
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        }
      },
      "cloud": {
        "iops_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "latency_raw": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "flexcache_raw": {
        "cache_miss_blocks": 10,
        "client_requested_blocks": 500,
        "status": "ok",
        "timestamp": "2017-01-25 11:20:13 +0000"
      },
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "nfs_ops_raw": {
        "access": {
          "count": 1000,
          "total_time": 200
        },
```

```
          "audit": {
            "count": 1000,
            "total_time": 200
          },
          "create": {
            "dir": {
              "count": 1000,
              "total_time": 200
            },
            "file": {
              "count": 1000,
              "total_time": 200
            },
            "other": {
              "count": 1000,
              "total_time": 200
            },
            "symlink": {
              "count": 1000,
              "total_time": 200
            }
          },
          "getattr": {
            "count": 1000,
            "total_time": 200
          },
          "link": {
            "count": 1000,
            "total_time": 200
          },
          "lock": {
            "count": 1000,
            "total_time": 200
          },
          "lookup": {
            "count": 1000,
            "total_time": 200
          },
          "open": {
            "count": 1000,
            "total_time": 200
          },
          "read": {
            "count": 1000,
            "total_time": 200,
            "volume_protocol_latency_histogram_counts": [
```

```
            0,
            0,
            0,
            0,
            0,
            15,
            35,
            100,
            200,
            200,
            300,
            500,
            500,
            500,
            1000,
            1000,
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
```

```
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
```

```
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        },
        "readdir": {
          "count": 1000,
          "total_time": 200
        },
        "readlink": {
          "count": 1000,
          "total_time": 200
        },
        "rename": {
          "count": 1000,
          "total_time": 200
        },
        "setattr": {
          "count": 1000,
          "total_time": 200
        },
        "unlink": {
          "count": 1000,
          "total_time": 200
        },
```

```
            "watch": {
              "count": 1000,
              "total_time": 200
            },
            "write": {
              "count": 1000,
              "total_time": 200,
              "volume_protocol_latency_histogram_counts": [
                0,
                0,
                0,
                0,
                0,
                15,
                35,
                100,
                200,
                200,
                300,
                500,
                500,
                500,
                1000,
                1000,
                800,
                500,
                500,
                300,
                200,
                50,
                40,
                15,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
                0,
```

```
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
```

```
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        }
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "status": [
      "string"
    ],
```

```
      "style": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tiering": {
        "object_tags": [
          "string"
        ],
        "policy": "string",
        "supported": true
      },
      "type": "string",
      "use_mirrored_aggregates": true,
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
|---|---|---|
| labels | array[string] | Labels for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

| Name | Type | Description |
|---|---|---|
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
|---|---|---|
| labels | array[string] | Labels for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification.

| Name | Type | Description |
|---|---|---|
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

analytics

Additional file system analytics information that is non-modifiable amongst all elements in the collection.

This property is only populated if file system analytics is enabled on the containing volume.

This analytics object captures properties that are non-modifiable amongst all elements included in the `records` array. The non-modifiable properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large.

| Name | Type | Description |
|---|---|---|
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

notices

| Name | Type | Description |
|---|---|---|
| code | string | An error code related to activity tracking. |
| message | string | A notice message related to activity tracking. |

unsupported_reason

| Name | Type | Description |
|---|---|---|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|---|---|---|
| notices | array[notices] | List of activity-tracking notices. |
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.<br>• enum: ["off", "on"]<br>• Introduced in: 9.10<br>• x-nullable: true |
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

initialization

with any other time period. This usually occurs when the data was at some point associated with a

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the analytics file system scan. |

scan_throttle_reason

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[string] | Arguments present in the warning message encountered. |
| code | string | Warning code indicating why scanner throttling is reported. |
| message | string | A message that provides details for scan throttling. |

unsupported_reason

| Name | Type | Description |
| --- | --- | --- |
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |

| Name | Type | Description |
| --- | --- | --- |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk. |
| file_count | integer | Number of descendants. |
| files_scanned | integer | Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| initialization | initialization | |
| report_time | string | Time of data collection. |
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| scan_throttle_reason | scan_throttle_reason | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume.<br>• enum: ["unknown", "initializing", "initialization_paused", "off", "on"]<br>• Introduced in: 9.8<br>• x-nullable: true |
| subdir_count | integer | Number of sub-directories. |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| total_files | integer | Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`. |
| unsupported_reason | unsupported_reason | |

anti_ransomware_volume_attack_detection_parameters

| Name | Type | Description |
|---|---|---|
| based_on_file_create_op_rate | boolean | Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_delete_op_rate | boolean | Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_rename_op_rate | boolean | Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes. |
| based_on_high_entropy_data_rate | boolean | Specifies whether a high entropy data rate should be considered for attack detection. |
| based_on_never_seen_before_file_extension | boolean | Specifies whether file extensions never seen before should be considered for attack detection. |
| block_device_auto_learned_encryption_threshold | integer | Specifies the block device auto learned encryption threshold. |
| file_create_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file create rate up to which it is considered normal behavior. |
| file_delete_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior. |
| file_rename_op_rate_surge_notify_percent | integer | Specifies the percent of surge in the file rename rate up to which it is considered normal behavior. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_surge_notify_percent | integer | Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on. |
| never_seen_before_file_extension_count_notify_threshold | integer | Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior. |
| never_seen_before_file_extension_duration_in_hours | integer | Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior. |
| relaxing_popular_file_extensions | boolean | Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on. |

_links

| Name | Type | Description |
| --- | --- | --- |
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

clear_suspect

Clear suspect status.

| Name | Type | Description |
|---|---|---|
| phase | string | Clear suspect phase. |
| start_time | string | Clear suspect start time. |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

space

| Name | Type | Description |
|---|---|---|
| snapshot_count | integer | Total number of Anti-ransomware backup snapshots. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup snapshots. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |

| Name | Type | Description |
|---|---|---|
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

suspect_files

| Name | Type | Description |
|---|---|---|
| count | integer | Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume. |
| entropy | string | Indicates the entropy level of this file type. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

historical_statistics

Typical usage values of volume workload.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

newly_observed_file_extensions

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Count of newly observed file extensions. |
| name | string | Name of the newly observed file extension. |

surge_statistics

Usage values of the volume's workload during surge.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

anti_ransomware_volume_workload

| Name | Type | Description |
| --- | --- | --- |
| file_extension_types_count | integer | Count of types of file extensions observed in the volume. |
| file_extensions_observed | array[string] | File extensions observed in the volume. |
| historical_statistics | historical_statistics | Typical usage values of volume workload. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| surge_statistics | surge_statistics | Usage values of the volume's workload during surge. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
| --- | --- | --- |
| attack_detected_by | string | This field specifies whether the attack was reported by `File Analysis` or `Encrypted data percentage analysis`. |
| attack_detection_parameters | anti_ransomware_volume_attack_detection_parameters | |
| attack_probability | string | Probability of a ransomware attack.<br>`none` No suspected ransomware activity.<br>`low` Minimally suspected ransomware activity.<br>`moderate` Moderately suspected ransomware activity.<br>`high` Significantly suspected ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |
| block_device_detection_start_time | string | This field specifies the block device evaluation start time. |
| block_device_detection_state | string | This field specifies the block device attack detection status.<br>`evaluation_period` Attack detection is currently in its evaluation phase.<br>`active_unsuitable_workload` Attack detection is active, but the current workload is not suitable for Anti-ransomware protection.<br>`active_suitable_workload` Attack detection is active, and the current workload is appropriate for Anti-ransomware protection. |
| clear_suspect | clear_suspect | Clear suspect status. |
| dry_run_start_time | string | Time when anti-ransomware monitoring `state` is set to dry-run value for starting evaluation mode.This field will no longer be supported in a future release. |
| event_log | event_log | |

| Name | Type | Description |
| --- | --- | --- |
| space | space | |
| state | string | Anti-ransomware state.<br>`disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.<br>`disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.<br>`dry_run` Anti-ransomware monitoring is enabled in the evaluation mode.<br>`enabled` Anti-ransomware monitoring is active on the volume.<br>`paused` Anti-ransomware monitoring is paused on the volume.<br>`enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.<br>`dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation.<br>For POST, the valid Anti-ransomware states are only `disabled`, `enabled` and `dry_run`, whereas for PATCH, `paused` is also valid along with the three valid states for POST. |
| surge_as_normal | boolean | Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| suspect_files | array[suspect_files] | |

| Name | Type | Description |
| --- | --- | --- |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |
| update_baseline_from_surge | boolean | Sets the observed surge value as the new baseline on a volume. |
| workload | anti_ransomware_volume_workload | |

application

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether asynchronous directory delete from the client is enabled on the volume. |
| trash_bin | string | Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used. |

autosize

| Name | Type | Description |
| --- | --- | --- |
| grow_threshold | integer | Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size.. |
| maximum | integer | Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume. |
| minimum | integer | Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume. |
| mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| shrink_threshold | integer | Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

parent_volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

clone

| Name | Type | Description |
| --- | --- | --- |
| has_flexclone | boolean | Specifies whether this volume is a parent of any FlexClone volume. |
| inherited_physical_used | integer | Inherited physical used from the clone's base snapshot. |
| inherited_savings | integer | Inherited savings from the clone's base snapshot. |
| is_flexclone | boolean | Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST. |
| lun_name | string | This optional parameter specifies the name of a LUN that will be non-disruptively migrated to the newly created FlexClone volume. If not specified, no LUNs are migrated. |
| parent_snapshot | snapshot_reference | |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| parent_volume | parent_volume | |

| Name | Type | Description |
| --- | --- | --- |
| qtree_name | string | This optional parameter specifies the name of the qtree containing the LUN that will be non-disruptively migrated to the newly created FlexClone volume. |
| split_complete_percent | integer | Percentage of FlexClone volume blocks split from its parent volume. |
| split_estimate | integer | Space required by the containing-aggregate to split the FlexClone volume. |
| split_initiated | boolean | This field is set when a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. Valid in PATCH. |

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the aggregate hosting the FlexGroup volume constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the FlexGroup volume constituent. |
| uuid | string | List of the node UUIDs hosting the FlexGroup volume constituent. |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |

| Name | Type | Description |
| --- | --- | --- |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |

| Name | Type | Description |
|---|---|---|
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| large_size_enabled | boolean | Specifies whether the support for large volumes and large files is enabled on the volume. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
| --- | --- | --- |
| aggregates | aggregates | |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| name | string | FlexGroup volume constituent name. |
| node | node | |
| space | space | |

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms. |

| Name | Type | Description |
| --- | --- | --- |
| inactive_days | integer | Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'. |
| mode | string | Specifies the mode of inactive data compression scanner. Valid for PATCH and GET. |
| operation_state | string | Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active". |
| status | string | Status of last inactive data compression scan on the volume. |
| threshold_inactive_time | string | Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num> |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform. |

scanner

| Name | Type | Description |
| --- | --- | --- |
| compression | boolean | Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms. |
| dedupe | boolean | Start deduplication if scanning old data. Valid for PATCH and GET. |
| scan_old_data | boolean | Indicates whether or not to scan old data. Valid for PATCH and GET. |
| state | string | State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active". |

space_savings

| Name | Type | Description |
| --- | --- | --- |
| compression | integer | Total disk space that is saved by compressing blocks on the referenced file system, in bytes. |
| compression_percent | integer | Percentage of total disk space that is saved by compressing blocks on the referenced file system. |
| dedupe | integer | Total disk space that is saved by deduplication and file cloning, in bytes. |
| dedupe_percent | integer | Percentage of total disk space that is saved by deduplication and file cloning. |
| dedupe_sharing | integer | Total disk space that is shared due to deduplication and file cloning. |
| total | integer | Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| total_percent | integer | Percentage of total disk space saved in the volume due to deduplication, compression and file cloning. |

efficiency

| Name | Type | Description |
| --- | --- | --- |
| application_io_size | string | Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms. |
| auto_state | string | Automatic deduplication schedule volume state.<br> auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically.<br> deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically. |
| compaction | string | The system can be enabled/disabled compaction.<br> inline ‐ Data will be compacted first and written to the volume.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled. |

| Name | Type | Description |
| --- | --- | --- |
| compression | string | The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be compressed first and written to the volume.<br> background ‐ Data will be written to the volume and compressed later.<br> both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled.<br><br>that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume. |
| compression_type | string | Compression type to use by compression. Valid for PATCH and GET. |

| Name | Type | Description |
| --- | --- | --- |
| cross_volume_dedupe | string | The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be cross volume deduped first and written to the volume.<br> background ‐ Data will be written to the volume and cross volume deduped later.<br> both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled. |
| dedupe | string | The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be deduped first and written to the volume.<br> background ‐ Data will be written to the volume and deduped later.<br> both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled. |
| has_savings | boolean | When true, indicates that the volume contains shared(deduplication, file clones) or compressed data. |

| Name | Type | Description |
| --- | --- | --- |
| idcs_scanner | idcs_scanner | Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| logging_enabled | boolean | When true, indicates that space savings for any newly-written data are being logged. |
| op_state | string | Sis status of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| ratio | number | Storage efficiency that does not include the savings provided by snapshots. |
| scanner | scanner | |
| schedule | string | Schedule associated with volume. |
| space_savings | space_savings | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP and Cloud Volumes ONTAP.<br> disabled ‐ All storage efficiency features are disabled.<br> mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others.<br> On FSx for ONTAP and Cloud Volumes ONTAP ‐<br>   enabled ‐ All supported storage efficiency features for the volume are enabled.<br>   custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled.<br> For other platforms ‐<br>   enabled ‐ At least one storage efficiency feature for the volume is enabled.<br><br>• enum: ["disabled", "enabled", "mixed", "custom"]<br>• Introduced in: 9.9<br>• x-nullable: true |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes. |
| type | string | Sis Type of the volume. |
| volume_path | string | Absolute volume path of the volume. |

status

| Name | Type | Description |
| --- | --- | --- |
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| action | string | This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. The following actions are supported for this field:<br> conversion_pause ‐ Pause an encryption conversion operation currently in progress.<br> conversion_resume ‐ Resume a paused encryption conversion operation.<br> rekey_pause ‐ Pause an encryption rekey operation currently in progress.<br> rekey_resume ‐ Resume a paused encryption rekey operation. |
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_create_time | string | Encryption key creation time of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state.<br> encrypted ‐ The volume is completely encrypted.<br> encrypting ‐ Encryption operation is in progress.<br> partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume.<br> rekeying. Encryption of volume with a new key is in progress.<br> unencrypted ‐ The volume is a plain-text one. |
| status | status | |
| type | string | Volume encryption type.<br> none ‐ The volume is a plain-text one.<br> volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption).<br> aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
| --- | --- | --- |
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |

| Name | Type | Description |
| --- | --- | --- |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies.<br> true ‐ File system is inconsistent.<br> false ‐ File system in not inconsistent. |

files

| Name | Type | Description |
| --- | --- | --- |
| inodefile_capacity | integer | Number of inodes that can currently be stored on the volume for user-visible files. This number dynamically increases as more user-visible files are created. |
| maximum | integer | The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH. |
| used | integer | Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online. |

flash_pool

| Name | Type | Description |
| --- | --- | --- |
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
| --- | --- | --- |
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes

hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| capacity_tier_optimized | boolean | Specifies whether this move should perform optimized volume move for the capacity tier, in which data in the capacity tier does not need to be copied over. |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | Export Policy |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ Unix-style security. |

| Name | Type | Description |
| --- | --- | --- |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth dight), first digit is assumed to be zero. |

nodes

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the volume. |
| uuid | string | List of the node UUIDs hosting the volume. |

policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed. |

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

last_error

Error information for the last failed file move on the constituent.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code of the last file move error on the constituent. |

| Name | Type | Description |
| --- | --- | --- |
| destination | integer | DSID of the destination constituent of the last file move error on the constituent. |
| file_id | integer | File ID of the last file move error on the constituent. |
| time | string | Time of the last file move error on the constituent. |

movement

Properties on this constituent related to file movement.

| Name | Type | Description |
| --- | --- | --- |
| file_moves_started | integer | Number of file moves started on this constituent. |
| last_error | last_error | Error information for the last failed file move on the constituent. |
| most_recent_start_time | string | Start time of the most recent file move on the constituent. |

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files. |

| Name | Type | Description |
| --- | --- | --- |
| footprint_invalid | integer | Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints. |
| in_snapshot | integer | Number of blocks skipped by the scanner on this constituent because of files in snapshots. |
| incompatible | integer | Number of blocks skipped by the scanner on this constituent because of incompatible files. |
| metadata | integer | Number of blocks skipped by the scanner on this constituent because of metadata files. |
| on_demand_destination | integer | Number of blocks skipped by the scanner on this constituent because of on demand destination files. |
| other | integer | Number of blocks skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of blocks skipped by the scanner on this constituent because of remote caches. |
| too_large | integer | Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size. |
| too_small | integer | Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of blocks skipped by the scanner on this constituent because of files fenced for write operations. |

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file. |
| footprint_invalid | integer | Number of files skipped by the scanner on this constituent because their space footprints are invalid. |
| in_snapshot | integer | Number of files skipped by the scanner on this constituent because they are trapped in snapshots. |
| incompatible | integer | Number of files skipped by the scanner on this constituent because they are incompatible. |
| metadata | integer | Number of files skipped by the scanner on this constituent because they metadata files. |
| on_demand_destination | integer | Number of files skipped by the scanner on this constituent because they are on demand destinations. |
| other | integer | Number of files skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of files skipped by the scanner on this constituent because they are remote caches. |

| Name | Type | Description |
| --- | --- | --- |
| too_large | integer | Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size. |
| too_small | integer | Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of files skipped by the scanner on this constituent because they are fenced for write operations. |

scanner

Properties related to determining which files to move and where to move them to.

| Name | Type | Description |
| --- | --- | --- |
| blocks_scanned | integer | Number of blocks scanned on this constituent. |
| blocks_skipped | blocks_skipped | Number of blocks skipped by the scanner on this constituent due to various reasons. |
| files_scanned | integer | Number of files scanned on this constituent. |
| files_skipped | files_skipped | Number of files skipped by the scanner on this constituent due to various reasons. |

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | Properties on this constituent related to file movement. |

| Name | Type | Description |
| --- | --- | --- |
| scanner | scanner | Properties related to determining which files to move and where to move them to. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| data_moved | integer | The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent. |
| engine | engine | Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions. |

| Name | Type | Description |
| --- | --- | --- |
| exclude_snapshots | boolean | Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value. |
| imbalance_percent | integer | Represents the percentage the volume is out of balance. |
| imbalance_size | integer | Represents how much the volume is out of balance, in bytes. |
| max_constituent_imbalance_percent | integer | Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active. |
| max_file_moves | integer | Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value. |

| Name | Type | Description |
| --- | --- | --- |
| max_runtime | string | This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num> |
| max_threshold | integer | Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value. |

| Name | Type | Description |
| --- | --- | --- |
| min_file_size | integer | Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance. |
| min_threshold | integer | Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value. |
| notices | array[error] | Capacity rebalancing notice messages. |
| runtime | string | Duration the capacity rebalancing operation has been running. |

| Name | Type | Description |
| --- | --- | --- |
| start_time | string | Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation.<br><br>While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither.<br><br>If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume.<br><br>If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property.<br><br>The following values apply to FlexGroup volumes.<br> not_running ‐ capacity rebalancing is not running on the volume.<br> starting ‐ used in a PATCH operation to start a capacity rebalancing operation.<br> rebalancing ‐ capacity rebalancing is running on the volume.<br> paused ‐ volume capacity rebalancing is paused on the volume.<br> stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation.<br> unknown ‐ the system was unable to determine the rebalancing state for the volume.<br><br>The following values apply to |

| Name | Type | Description |
| --- | --- | --- |
| stop_time | string | Time when the capacity rebalancing operation stopped. |
| target_used | integer | Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents. |
| used_for_imbalance | integer | Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size. |

retention

the system was unable to determine the rebalancing state for the constituent.

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num> |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |

| Name | Type | Description |
| --- | --- | --- |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | retention | |
| type | string | The SnapLock type of the volume.<br> compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files.<br> enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume.<br> non_snaplock ‐ Indicates the volume is non-snaplock. |
| unspecified_retention_file_count | integer | Indicates the number of files with an unspecified retention time in the volume. |

destinations

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination.<br>• readOnly: 1<br>• Introduced in: 9.9<br>• x-nullable: true |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
| --- | --- | --- |
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |

| Name | Type | Description |
| --- | --- | --- |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |
| used_by_snapshots | integer | Size that is logically used across all snapshots in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

autodelete

| Name | Type | Description |
| --- | --- | --- |
| commitment | string | By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try. |
| defer_delete | string | Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created. |

| Name | Type | Description |
| --- | --- | --- |
| delete_order | string | Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first. |
| enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| prefix | string | Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value. |
| target_free_space | integer | Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized. |
| trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete | autodelete | |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead. |

| Name | Type | Description |
| --- | --- | --- |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead. |
| reserve_available | integer | Size available for snapshots within the snapshot reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for snapshot usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| auto_adaptive_compression_footprint_data_reduction | integer | Savings achieved due to Auto Adaptive Compression, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |

| Name | Type | Description |
| --- | --- | --- |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| capacity_tier_footprint_data_reduction | integer | Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes. |
| compaction_footprint_data_reduction | integer | Savings achieved due to Data Compaction, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| effective_total_footprint | integer | Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| is_used_stale | boolean | Specifies if the virtual space used is stale. |
| large_size_enabled | boolean | Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB. |

| Name | Type | Description |
| --- | --- | --- |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. Physical used is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for snapshots from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for snapshot creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

access

Raw count and latency data for access operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

dir

Raw count and latency data for directory-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

file

Raw count and latency data for file-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

symlink

Raw count and latency data for symlink-create operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

create

Raw count and latency data for create operations.

| Name | Type | Description |
|---|---|---|
| dir | dir | Raw count and latency data for directory-create operations. |
| file | file | Raw count and latency data for file-create operations. |
| other | other | Raw count and latency data for create operations on objects other than files, directories and symlinks. |
| symlink | symlink | Raw count and latency data for symlink-create operations. |

getattr

Raw count and latency data for getattr operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

link

Raw count and latency data for link operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lock

Raw count and latency data for lock operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lookup

Raw count and latency data for lookup operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

open

Raw count and latency data for open operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

readdir

Raw count and latency data for readdir operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

readlink

Raw count and latency data for readlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

rename

Raw count and latency data for rename operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

setattr

Raw count and latency data for setattr operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

unlink

Raw count and latency data for unlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

| Name | Type | Description |
| --- | --- | --- |
| access | access | Raw count and latency data for access operations. |
| audit | audit | Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations. |
| create | create | Raw count and latency data for create operations. |
| getattr | getattr | Raw count and latency data for getattr operations. |
| link | link | Raw count and latency data for link operations. |
| lock | lock | Raw count and latency data for lock operations. |
| lookup | lookup | Raw count and latency data for lookup operations. |
| open | open | Raw count and latency data for open operations. |
| read | read | Raw count and latency data for read operations, including histograms categorizing operations by size and latency. |
| readdir | readdir | Raw count and latency data for readdir operations. |
| readlink | readlink | Raw count and latency data for readlink operations. |
| rename | rename | Raw count and latency data for rename operations. |
| setattr | setattr | Raw count and latency data for setattr operations. |
| unlink | unlink | Raw count and latency data for unlink operations. |

| Name | Type | Description |
| --- | --- | --- |
| watch | watch | Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations. |
| write | write | Raw count and latency data for write operations, including histograms categorizing operations by size and latency. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

| Name | Type | Description |
| --- | --- | --- |
| access | access | Raw count and latency data for access operations. |
| audit | audit | Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations. |
| create | create | Raw count and latency data for create operations. |
| getattr | getattr | Raw count and latency data for getattr operations. |

| Name | Type | Description |
| --- | --- | --- |
| link | link | Raw count and latency data for link operations. |
| lock | lock | Raw count and latency data for lock operations. |
| lookup | lookup | Raw count and latency data for lookup operations. |
| open | open | Raw count and latency data for open operations. |
| read | read | Raw count and latency data for read operations, including histograms categorizing operations by size and latency. |
| readdir | readdir | Raw count and latency data for readdir operations. |
| readlink | readlink | Raw count and latency data for readlink operations. |
| rename | rename | Raw count and latency data for rename operations. |
| setattr | setattr | Raw count and latency data for setattr operations. |
| unlink | unlink | Raw count and latency data for unlink operations. |
| watch | watch | Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations. |
| write | write | Raw count and latency data for write operations, including histograms categorizing operations by size and latency. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| cifs_ops_raw | cifs_ops_raw | Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| nfs_ops_raw | nfs_ops_raw | Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM containing the volume. Required on POST.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tiering

| Name | Type | Description |
| --- | --- | --- |
| min_cooling_days | integer | This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy. |
| object_tags | array[string] | This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value". |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br> all ‐ This policy allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br> auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store<br> none ‐ Volume blocks will not be tiered to the cloud store.<br> snapshot_only ‐ This policy allows tiering of only the volume snapshots not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |

| Name | Type | Description |
| --- | --- | --- |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |

| Name | Type | Description |
| --- | --- | --- |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |

| Name | Type | Description |
| --- | --- | --- |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| metric | metric | Performance numbers, such as IOPS, latency and throughput. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | Type of the volume.<br> rw ‐ read-write volume.<br> dp ‐ data-protection volume.<br> ls ‐ load-sharing <code>dp</code> volume. Valid in GET. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a volume on an SVM and storage aggregates**

POST `/storage/volumes`

**Introduced In:** 9.6

Creates a volume on a specified SVM and storage aggregates.

**Required properties**

- `svm.uuid` or `svm.name` - Existing SVM in which to create the volume.
- `name` - Name of the volume.
- `aggregates.name` or `aggregates.uuid` - Existing aggregates in which to create the volume.

#### Default property values

- `state` - *online*
- `size` - *20MB*
- `style` - *flexvol*
- `type` - *rw*
- `encryption.enabled` - *false*
- `snapshot_policy.name` - *default*
- `guarantee.type` - *volume*
- `anti_ransomware.state` - *default*

#### Related ONTAP commands

- `volume create`
- `volume clone create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| activity_tracking | activity_tracking | |

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier.<br>This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| constituents_per_aggregate | integer | Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |

| Name | Type | Description |
| --- | --- | --- |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
| --- | --- | --- |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| optimize_aggregates | boolean | Specifies whether to create the constituents of the FlexGroup volume on the aggregates specified in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is 'true', the system optimizes the ordering of the aggregates specified. If this value is false, the order of the aggregates is unchanged. The default value is 'false'. |
| qos | qos | QoS information |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |

| Name | Type | Description |
| --- | --- | --- |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | Type of the volume.<br> rw ‐ read-write volume.<br> dp ‐ data-protection volume.<br> ls ‐ load-sharing <code>dp</code> volume. Valid in GET. |

| Name | Type | Description |
| --- | --- | --- |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

**Example request**


```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "access_time_enabled": true,
  "activity_tracking": {
    "notices": [
      {
        "code": "124518424",
        "message": "The top metrics report contains partial data for
read operations because NFSv4 reads using Multi-Processor I/O (MPIO)
are not tracked."
      }
    ],
    "state": "string",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
```

```
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
```

```
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "bytes_used": 15436648448,
    "file_count": 21134,
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "report_time": "2024-11-06 13:57:15 -0500",
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
    },
    "state": "string",
    "subdir_count": 35,
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_detected_by": "string",
    "attack_detection_parameters": {
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "time": "2021-06-01 20:36:41 +0530"
      }
    ],
    "block_device_detection_start_time": "string",
```

```
    "block_device_detection_state": "string",
    "clear_suspect": {
      "phase": "string",
      "start_time": "string"
    },
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 23:16:20 +0530"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
      ],
      "historical_statistics": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
```

```
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "surge_statistics": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 23:16:20 +0530"
      },
      "surge_usage": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
            "name": [
              "lockile"
            ]
          }
        ],
        "time": "2021-12-01 23:16:20 +0530"
      },
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
```

```
      }
    },
    "application": {
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
    },
    "asynchronous_directory_delete": {
      "trash_bin": "string"
    },
    "autosize": {
      "mode": "string"
    },
    "clone": {
      "inherited_physical_used": 0,
      "inherited_savings": 0,
      "lun_name": "string",
      "parent_snapshot": {
        "name": "this_snapshot",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "parent_svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "parent_volume": {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "qtree_name": "string",
      "split_complete_percent": 0,
      "split_estimate": 0,
      "split_initiated": true
    },
    "cloud_retrieval_policy": "string",
    "comment": "string",
    "consistency_group": {
      "name": "consistency_group_1",
      "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
    },
    "constituents": [
      {
        "aggregates": {
          "name": "string",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        },
        "movement": {
```

```
        "cutover_window": 30,
        "destination_aggregate": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "state": "replicating",
        "tiering_policy": "string"
      },
      "name": "string",
      "node": {
        "name": "string",
        "uuid": "string"
      },
      "space": {
        "available": 0,
        "block_storage_inactive_user_data": 0,
        "capacity_tier_footprint": 0,
        "footprint": 0,
        "local_tier_footprint": 0,
        "logical_space": {
          "available": 0,
          "used_by_afs": 0
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 19:00:00 +0000",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
```

```
    "dedupe": "string",
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
    "policy": {
      "name": "string"
    },
    "progress": "string",
    "ratio": 0,
    "scanner": {
      "compression": true,
      "dedupe": true,
      "scan_old_data": true
    },
    "schedule": "string",
    "space_savings": {
      "compression": 0,
      "compression_percent": 0,
      "dedupe": 0,
      "dedupe_percent": 0,
      "dedupe_sharing": 0,
      "total": 0,
      "total_percent": 0
    },
    "state": "string",
    "storage_efficiency_mode": "string",
    "type": "string",
    "volume_path": "string"
  },
  "encryption": {
    "action": "string",
    "key_create_time": "2022-01-01 19:00:00 +0000",
    "key_id": "string",
    "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
    "rekey": true,
```

```
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "inodefile_capacity": 0,
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "string",
    "cache_retention_priority": "string",
    "caching_policy": "string"
  },
  "flexcache_endpoint_type": "string",
  "flexgroup": {
    "name": "my_flexgroup",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "granular_data_mode": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "movement": {
    "capacity_tier_optimized": true,
    "percent_complete": 0,
    "start_time": "2020-12-07 03:45:12 -0500"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 493
  },
  "nodes": [
```

```
      {
        "name": "string",
        "uuid": "string"
      }
    ],
    "qos": {
      "policy": {
        "max_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput": [
          "900KB/s",
          "500MB/s",
          "120GB/s",
          "5000IOPS",
          "5000IOPS,500KB/s",
          "2500IOPS,100MB/s",
          "1000IOPS,25MB/s"
        ],
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "queue_for_encryption": true,
    "quota": {
      "state": "string"
    },
    "rebalancing": {
      "data_moved": 0,
      "engine": {
        "movement": {
          "file_moves_started": 0,
          "last_error": {
            "code": 0,
            "destination": 0,
            "file_id": 0,
```

```
          "time": "2018-06-04 19:00:00 +0000"
        },
        "most_recent_start_time": "2018-06-04 19:00:00 +0000"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      }
    },
    "exclude_snapshots": true,
    "imbalance_percent": 0,
    "imbalance_size": 0,
    "max_constituent_imbalance_percent": 0,
    "notices": [
      {
        "arguments": [
```

```
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "sizing_method": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 19:00:00 +0000",
    "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
    "is_audit_log": 1,
    "litigation_count": 10,
    "privileged_delete": "enabled",
    "retention": {
      "default": "P30Y",
      "maximum": "P30Y",
      "minimum": "P30Y"
    },
    "type": "enterprise",
    "unspecified_retention_file_count": 10
  },
  "snapshot_count": 0,
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 0,
    "cross_volume_dedupe_metafiles_footprint": 0,
```

```
      "cross_volume_dedupe_metafiles_temporary_footprint": 0,
      "dedupe_metafiles_footprint": 0,
      "dedupe_metafiles_temporary_footprint": 0,
      "delayed_free_footprint": 0,
      "effective_total_footprint": 0,
      "file_operation_metadata": 0,
      "filesystem_size": 0,
      "footprint": 0,
      "local_tier_footprint": 0,
      "logical_space": {
        "available": 0,
        "used": 0,
        "used_by_afs": 0,
        "used_by_snapshots": 0,
        "used_percent": 0
      },
      "max_size": "string",
      "metadata": 0,
      "over_provisioned": 0,
      "overwrite_reserve": 0,
      "overwrite_reserve_used": 0,
      "percent_used": 0,
      "performance_tier_footprint": 0,
      "size_available_for_snapshots": 0,
      "snapmirror_destination_footprint": 0,
      "snapshot": {
        "reserve_available": 0,
        "reserve_size": 0,
        "space_used_percent": 0,
        "used": 0
      },
      "snapshot_reserve_unusable": 0,
      "snapshot_spill": 0,
      "total_footprint": 0,
      "total_metadata": 0,
      "total_metadata_footprint": 0,
      "used": 0,
      "user_data": 0,
      "volume_guarantee_footprint": 0
    },
    "state": "string",
    "status": [
      "string"
    ],
    "style": "string",
    "svm": {
```

```
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string"
  },
  "type": "string",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 460770 | [Job "jobid"] Job failed. |
| 524561 | Unable to get the information for the specified SVM's root volume because of the specified reason. |
| 524597 | Failed to create the clone volume on the node. |
| 524597 | Failed to create clone volume on the node. |
| 524601 | Parent volume "name" not found. |
| 524819 | Failed to create the clone volume. |
| 787140 | One of "aggregates.uuid", "aggregates.name", or "style" must be provided. |
| 787141 | The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates. |
| 917508 | The target aggregate could not be found. |
| 917526 | The volume name specified is a duplicate. |
| 917534 | The value specified for field "-size" is too small. Update the field "-size" with the minimum size allowed and retry. |
| 917536 | Job completed with specified error. |
| 917551 | Specified snapshot policy not found. |
| 917597 | Provisioning a volume on a root aggregate is not supported. |
| 917775 | Volume \"{0}:{1}\" is not online. |
| 917829 | Volume autosize grow threshold must be larger than autosize shrink threshold. |
| 917831 | Volume minimum autosize must be smaller than the maximum autosize. |
| 917835 | Maximum allowed snapshot.reserve_percent value during a volume creation is 90. Use PATCH to set it to a higher value after the volume has been created. |
| 917886 | Volume name is too long. It must be 203 characters or less. |
| 917887 | The first character of the volume name must be a letter or underscore. |

| Error Code | Description |
| --- | --- |
| 918157 | Node \"{0}\" in cluster \"{1}\" failed to allocate encryption resources. Wait a few minutes, and try the operation again. If the issue persists, reboot the node and try the operation again. |
| 918162 | Node \"{0}\" in cluster \"{1}\" is not reachable. Creation of the first encrypted volume requires encryption resources to be allocated on both node \"{2}\" and its HA partner node \"{0}\". Make sure that the node and its partner are healthy, and then try the operation again. |
| 918191 | Flexvol tiering min cooling days requires an effective cluster version of ONTAP 9.4 or later. |
| 918194 | Tiering min cooling days not supported for SVMDR. |
| 918195 | Tiering min cooling days not supported for non data volumes. |
| 918196 | Tiering min cooling days not allowed for the provided tiering policy. |
| 918215 | FlexGroup tiering min cooling days requires an effective cluster version of ONTAP 9.5 or later. |
| 918232 | The specified volume identifier fields must be provided. |
| 918233 | The target field cannot be specified for this operation. |
| 918236 | The specified "parent_volume.uuid" and "parent_volume.name" do not refer to the same volume. |
| 918240 | The target style is an invalid volume style. |
| 918241 | The target style is an unsupported volume style for volume creation. |
| 918242 | When creating a flexible volume, exactly one aggregate must be specified via either "aggregates.name" or "aggregates.uuid". |
| 918243 | The specified snapshot UUID is not correct for the specified snapshot name. |
| 918244 | Invalid "volume.type" for clone volume. |
| 918247 | Specifying a value is not valid for a volume FlexClone creation. |
| 918252 | "nas.path" is invalid. |
| 918271 | Failed to get valid export policy. |
| 918290 | cloud retrieval policy requires an effective cluster version of 9.8 or later. |

| Error Code | Description |
| --- | --- |
| 918291 | Invalid volume cloud retrieval policy for the provided tiering policy. |
| 918292 | cloud retrieval policy not supported for non data volume. |
| 918334 | Cache cloud retrieval policy requires an effective cluster version of 9.14 or later. |
| 918521 | The volume maximum autosize must be smaller than or equal to the maximum volume size. |
| 918524 | Volume minimum autosize must be less than or equal to the current volume size. |
| 918652 | "error" is an invalid value for field "-state". Valid values are "online", "offline" and "restricted". |
| 918701 | The specified operation on the volume endpoint is not supported on this platform. |
| 1638400 | Failed to retrieve snapshot information. |
| 1638587 | The specified snapshot_policy.name and snapshot_policy.uuid refer to different snapshot policies. |
| 1638593 | Operation failed because "snapdb" is disabled. |
| 1638624 | Internal error. Failed to lookup snapshot tags for \"{0}\". |
| 2621462 | The target SVM does not exist. |
| 2621706 | The specified "svm.uuid" and "svm.name" do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied. |
| 6881654 | Disabling storage efficiency features is not allowed on this platform. |
| 13107245 | The specified FlexGroup volume layout exceeds the maximum number of volumes on node. |
| 13107307 | specifying this parameter when creating a flexible volume is not supported. |
| 13107326 | Request to provision FlexGroup volume "name" failed because the requested size is greater than the maximum size. |
| 13107341 | At least one valid aggregate assigned to SVM \"{0}\" with sufficient space and a homogeneous storage type is required on every node on the cluster to provision a FlexGroup volume. |
| 13107349 | Operation is only supported on flexible volumes and FlexGroup volumes. |

| Error Code | Description |
| --- | --- |
| 13107405 | Another volume is currently being created with the name "name" in SVM "svm.name". |
| 13107413 | Creating a FlexGroup volume is not supported on All SAN Arrays and systems that support large LUNs. |
| 13109258 | Cannot enable granular data on volume "name" in Vserver "svm.name". This setting can only be enabled on FlexGroups. |
| 13109260 | Failed to enable granular data on the volume. |
| 13109290 | Specifying the "aggregates" parameter on this platform requires that the "constituents_per_aggregate" parameter is also specified. Instead, consider using the "constituent_count" parameter to specify how many constituents to create, and the system will automatically select where to place them. |
| 13109291 | Specifying the "constituents_per_aggregate" parameter requires that the "aggregates" parameter is also specified. |
| 13565983 | A value of zero is not supported for \"-uid\" or \"-gid\". |
| 13566001 | LUN name length cannot be zero. |
| 13566003 | Specifying the "qtree_name" parameter requires that the "lun_name" parameter is also specified. |
| 65537463 | Volume encryption keys (VEK) cannot be created or deleted for data Vserver \"{0}\". External key management has been configured for data Vserver \"{0}\" but ONTAP is not able to encrypt or decrypt with the key manager. Resolve the external key manager key issues at the key manager’s portal before creating any new encrypted volumes. |
| 65537529 | Encrypted volumes cannot be created or deleted for Vserver \"{0}\" as a rekey operation for the vserver is in progress. Try creating the encrypted volume again after some time. If the problem persists, run the rekey operation again after some time. |
| 65537600 | Encrypted volumes cannot be created or deleted for Vserver \"{0}\" while the enabled keystore configuration is being switched. If a previous attempt to switch the keystore configuration failed, or was interrupted, the system will continue to prevent encrypted volume creation for Vserver \"{0}\". Use the \"security key-manager keystore enable\" command to re-run and complete the operation. |

| Error Code | Description |
| --- | --- |
| 65539430 | Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation. |
| 65539431 | Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being disabled. |
| 111411203 | Failed to retrieve the volume file system analytics report on the volume. |
| 111411204 | Internal error. Failed to retrieve the volume file system analytics report on the volume. |
| 111411205 | File system analytics requires an effective cluster version of 9.8 or later. |
| 111411206 | The specified "analytics.state" is invalid. |
| 111411207 | File system analytics cannot be enabled on volumes that contain LUNs. |
| 111411207 | Volume file system analytics is not supported on volumes that contain LUNs. |
| 111411209 | Volume file system analytics is not supported on FlexCache volumes. |
| 111411210 | Volume file system analytics is not supported on audit staging volumes. |
| 111411211 | Volume file system analytics is not supported on object store server volumes. |
| 111411212 | Volume file system analytics is not supported on SnapMirror destination volumes. |
| 111411215 | Internal error. Volume file system analytics report timed out for volume volume.name in SVM svm.name. |
| 111411216 | Enabling or disabling volume file system analytics is not supported on individual FlexGroup constituents. |
| 111411217 | Volume file system analytics is not supported on SnapLock volumes. |
| 111411230 | Volume file system analytics is not supported on volumes that contain NVMe namespaces. |
| 111411241 | Volume file system analytics is not supported on All SAN Array clusters. |
| 111411252 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent`" available space in the volume, then try the operation again. |

| Error Code | Description |
| --- | --- |
| 111411253 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent" available space in all constituents of the FlexGroup, then try the operation again." |
| 111411257 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Increase the volume size to "size", then try the operation again." |
| 111412203 | Volume file system analytics is not enabled on the volume. |
| 111415208 | Cannot create a clone of volume "name" in SVM "svm.name" because the file system analytics state is not set to "off" or "on". Use REST API GET method "/api/storage/volumes/uuid?fields=analytics.state" to check the analytics state. |
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 196608308 | Node \"{0}\" in cluster \"{1}\" is not reachable. Creation of the first NAE (NetApp Aggregate Encryption) aggregate requires encryption resources to be allocated on both node \"{2}\" and its HA partner node \"{0}\". Make sure that the node and its partner are healthy, and then try the operation again. |

| Error Code | Description |
| --- | --- |
| 196608399 | One or more nodes in cluster \"{0}\" are not reachable. Creation of the first NAE (NetApp Aggregate Encryption) aggregate requires encryption resources to be allocated on all nodes in DR group \"{1}\". Ensure that all the nodes are healthy and then try the operation again. |
| 196608400 | One or more nodes in cluster \"{0}\" are not reachable. Creation of the first encrypted volume requires encryption resources to be allocated on all nodes in DR group \"{1}\". Ensure that all the nodes are healthy and then try the operation again. |
| 196608401 | Node \"{0}\" in cluster \"{1}\" is not reachable. Creation of the first NAE (NetApp Aggregate Encryption) aggregate requires encryption resources to be allocated on node \"{0}\". Make sure that the node is healthy, and then try the operation again. |
| 196608402 | Node \"{0}\" in cluster \"{1}\" is not reachable. Creation of the first encrypted volume requires encryption resources to be allocated on node \"{0}\". Make sure that the node is healthy, and then try the operation again. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

notices

| Name | Type | Description |
|---|---|---|
| code | string | An error code related to activity tracking. |
| message | string | A notice message related to activity tracking. |

unsupported_reason

| Name | Type | Description |
|---|---|---|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|---|---|---|
| notices | array[notices] | List of activity-tracking notices. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.<br>• enum: ["off", "on"]<br>• Introduced in: 9.10<br>• x-nullable: true |
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

initialization

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the analytics file system scan. |

scan_throttle_reason

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[string] | Arguments present in the warning message encountered. |
| code | string | Warning code indicating why scanner throttling is reported. |
| message | string | A message that provides details for scan throttling. |

unsupported_reason

| Name | Type | Description |
| --- | --- | --- |
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk. |
| file_count | integer | Number of descendants. |
| files_scanned | integer | Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| initialization | initialization | |
| report_time | string | Time of data collection. |
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| scan_throttle_reason | scan_throttle_reason | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume.<br>• enum: ["unknown", "initializing", "initialization_paused", "off", "on"]<br>• Introduced in: 9.8<br>• x-nullable: true |
| subdir_count | integer | Number of sub-directories. |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| total_files | integer | Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`. |
| unsupported_reason | unsupported_reason | |

anti_ransomware_volume_attack_detection_parameters

| Name | Type | Description |
|---|---|---|
| based_on_file_create_op_rate | boolean | Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_delete_op_rate | boolean | Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_rename_op_rate | boolean | Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes. |
| based_on_high_entropy_data_rate | boolean | Specifies whether a high entropy data rate should be considered for attack detection. |
| based_on_never_seen_before_file_extension | boolean | Specifies whether file extensions never seen before should be considered for attack detection. |
| block_device_auto_learned_encryption_threshold | integer | Specifies the block device auto learned encryption threshold. |
| file_create_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file create rate up to which it is considered normal behavior. |
| file_delete_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior. |
| file_rename_op_rate_surge_notify_percent | integer | Specifies the percent of surge in the file rename rate up to which it is considered normal behavior. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_surge_notify_percent | integer | Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on. |
| never_seen_before_file_extension_count_notify_threshold | integer | Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior. |
| never_seen_before_file_extension_duration_in_hours | integer | Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior. |
| relaxing_popular_file_extensions | boolean | Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on. |

anti_ransomware_attack_report

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

clear_suspect

Clear suspect status.

| Name | Type | Description |
|---|---|---|
| phase | string | Clear suspect phase. |
| start_time | string | Clear suspect start time. |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

space

| Name | Type | Description |
|---|---|---|
| snapshot_count | integer | Total number of Anti-ransomware backup snapshots. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup snapshots. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

suspect_files

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume. |
| entropy | string | Indicates the entropy level of this file type. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

historical_statistics

Typical usage values of volume workload.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

newly_observed_file_extensions

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Count of newly observed file extensions. |
| name | string | Name of the newly observed file extension. |

surge_statistics

Usage values of the volume's workload during surge.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

anti_ransomware_volume_workload

| Name | Type | Description |
| --- | --- | --- |
| file_extension_types_count | integer | Count of types of file extensions observed in the volume. |
| file_extensions_observed | array[string] | File extensions observed in the volume. |
| historical_statistics | historical_statistics | Typical usage values of volume workload. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| surge_statistics | surge_statistics | Usage values of the volume's workload during surge. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
| --- | --- | --- |
| attack_detected_by | string | This field specifies whether the attack was reported by `File Analysis` or `Encrypted data percentage analysis`. |
| attack_detection_parameters | anti_ransomware_volume_attack_detection_parameters | |
| attack_probability | string | Probability of a ransomware attack.<br>`none` No suspected ransomware activity.<br>`low` Minimally suspected ransomware activity.<br>`moderate` Moderately suspected ransomware activity.<br>`high` Significantly suspected ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |
| block_device_detection_start_time | string | This field specifies the block device evaluation start time. |
| block_device_detection_state | string | This field specifies the block device attack detection status.<br>`evaluation_period` Attack detection is currently in its evaluation phase.<br>`active_unsuitable_workload` Attack detection is active, but the current workload is not suitable for Anti-ransomware protection.<br>`active_suitable_workload` Attack detection is active, and the current workload is appropriate for Anti-ransomware protection. |
| clear_suspect | clear_suspect | Clear suspect status. |
| dry_run_start_time | string | Time when anti-ransomware monitoring `state` is set to dry-run value for starting evaluation mode.This field will no longer be supported in a future release. |
| event_log | event_log | |

| Name | Type | Description |
| --- | --- | --- |
| space | space | |
| state | string | Anti-ransomware state.<br>`disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.<br>`disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.<br>`dry_run` Anti-ransomware monitoring is enabled in the evaluation mode.<br>`enabled` Anti-ransomware monitoring is active on the volume.<br>`paused` Anti-ransomware monitoring is paused on the volume.<br>`enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.<br>`dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation.<br>For POST, the valid Anti-ransomware states are only `disabled`, `enabled` and `dry_run`, whereas for PATCH, `paused` is also valid along with the three valid states for POST. |
| surge_as_normal | boolean | Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| suspect_files | array[suspect_files] | |

| Name | Type | Description |
| --- | --- | --- |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |
| update_baseline_from_surge | boolean | Sets the observed surge value as the new baseline on a volume. |
| workload | anti_ransomware_volume_workload | |

application

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether asynchronous directory delete from the client is enabled on the volume. |
| trash_bin | string | Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used. |

autosize

| Name | Type | Description |
| --- | --- | --- |
| grow_threshold | integer | Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size.. |
| maximum | integer | Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume. |
| minimum | integer | Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume. |
| mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| shrink_threshold | integer | Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

parent_volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

clone

| Name | Type | Description |
| --- | --- | --- |
| has_flexclone | boolean | Specifies whether this volume is a parent of any FlexClone volume. |
| inherited_physical_used | integer | Inherited physical used from the clone's base snapshot. |
| inherited_savings | integer | Inherited savings from the clone's base snapshot. |
| is_flexclone | boolean | Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST. |
| lun_name | string | This optional parameter specifies the name of a LUN that will be non-disruptively migrated to the newly created FlexClone volume. If not specified, no LUNs are migrated. |
| parent_snapshot | snapshot_reference | |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| parent_volume | parent_volume | |

| Name | Type | Description |
| --- | --- | --- |
| qtree_name | string | This optional parameter specifies the name of the qtree containing the LUN that will be non-disruptively migrated to the newly created FlexClone volume. |
| split_complete_percent | integer | Percentage of FlexClone volume blocks split from its parent volume. |
| split_estimate | integer | Space required by the containing-aggregate to split the FlexClone volume. |

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the aggregate hosting the FlexGroup volume constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
|---|---|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the FlexGroup volume constituent. |
| uuid | string | List of the node UUIDs hosting the FlexGroup volume constituent. |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| large_size_enabled | boolean | Specifies whether the support for large volumes and large files is enabled on the volume. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
| --- | --- | --- |
| aggregates | aggregates | |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| name | string | FlexGroup volume constituent name. |
| node | node | |
| space | space | |

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms. |
| inactive_days | integer | Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'. |
| mode | string | Specifies the mode of inactive data compression scanner. Valid for PATCH and GET. |
| operation_state | string | Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active". |
| status | string | Status of last inactive data compression scan on the volume. |
| threshold_inactive_time | string | Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num> |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform. |

scanner

space_savings

| Name | Type | Description |
| --- | --- | --- |
| compression | integer | Total disk space that is saved by compressing blocks on the referenced file system, in bytes. |
| compression_percent | integer | Percentage of total disk space that is saved by compressing blocks on the referenced file system. |
| dedupe | integer | Total disk space that is saved by deduplication and file cloning, in bytes. |
| dedupe_percent | integer | Percentage of total disk space that is saved by deduplication and file cloning. |
| dedupe_sharing | integer | Total disk space that is shared due to deduplication and file cloning. |
| total | integer | Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes. |
| total_percent | integer | Percentage of total disk space saved in the volume due to deduplication, compression and file cloning. |

efficiency

| Name | Type | Description |
| --- | --- | --- |
| application_io_size | string | Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms. |
| auto_state | string | Automatic deduplication schedule volume state.<br> auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically.<br> deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically. |
| compaction | string | The system can be enabled/disabled compaction.<br> inline ‐ Data will be compacted first and written to the volume.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled. |

| Name | Type | Description |
| --- | --- | --- |
| compression | string | The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be compressed first and written to the volume.<br> background ‐ Data will be written to the volume and compressed later.<br> both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled.<br><br>that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume. |
| compression_type | string | Compression type to use by compression. Valid for PATCH and GET. |

| Name | Type | Description |
| --- | --- | --- |
| cross_volume_dedupe | string | The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be cross volume deduped first and written to the volume.<br> background ‐ Data will be written to the volume and cross volume deduped later.<br> both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled. |
| dedupe | string | The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be deduped first and written to the volume.<br> background ‐ Data will be written to the volume and deduped later.<br> both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled. |
| has_savings | boolean | When true, indicates that the volume contains shared(deduplication, file clones) or compressed data. |

| Name | Type | Description |
| --- | --- | --- |
| idcs_scanner | idcs_scanner | Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| logging_enabled | boolean | When true, indicates that space savings for any newly-written data are being logged. |
| op_state | string | Sis status of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| ratio | number | Storage efficiency that does not include the savings provided by snapshots. |
| scanner | scanner | |
| schedule | string | Schedule associated with volume. |
| space_savings | space_savings | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP and Cloud Volumes ONTAP.<br> disabled ‐ All storage efficiency features are disabled.<br> mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others.<br> On FSx for ONTAP and Cloud Volumes ONTAP ‐<br>   enabled ‐ All supported storage efficiency features for the volume are enabled.<br>   custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled.<br> For other platforms ‐<br>   enabled ‐ At least one storage efficiency feature for the volume is enabled.<br>• enum: ["disabled", "enabled", "mixed", "custom"]<br>• Introduced in: 9.9<br>• x-nullable: true |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes. |
| type | string | Sis Type of the volume. |
| volume_path | string | Absolute volume path of the volume. |

status

| Name | Type | Description |
| --- | --- | --- |
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| action | string | This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. The following actions are supported for this field:<br> conversion_pause ‐ Pause an encryption conversion operation currently in progress.<br> conversion_resume ‐ Resume a paused encryption conversion operation.<br> rekey_pause ‐ Pause an encryption rekey operation currently in progress.<br> rekey_resume ‐ Resume a paused encryption rekey operation. |
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_create_time | string | Encryption key creation time of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| key_manager_attribute | string | Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM. |
| state | string | Volume encryption state.<br> encrypted ‐ The volume is completely encrypted.<br> encrypting ‐ Encryption operation is in progress.<br> partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume.<br> rekeying. Encryption of volume with a new key is in progress.<br> unencrypted ‐ The volume is a plain-text one. |
| status | status | |
| type | string | Volume encryption type.<br> none ‐ The volume is a plain-text one.<br> volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption).<br> aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
| --- | --- | --- |
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies.<br> true ‐ File system is inconsistent.<br> false ‐ File system in not inconsistent. |

files

| Name | Type | Description |
| --- | --- | --- |
| inodefile_capacity | integer | Number of inodes that can currently be stored on the volume for user-visible files. This number dynamically increases as more user-visible files are created. |
| maximum | integer | The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH. |
| used | integer | Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online. |

flash_pool

| Name | Type | Description |
| --- | --- | --- |
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
| --- | --- | --- |
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes

hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |

export_policy

Export Policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | Export Policy |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |

| Name | Type | Description |
| --- | --- | --- |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ Unix-style security. |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth dight), first digit is assumed to be zero. |

nodes

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the volume. |
| uuid | string | List of the node UUIDs hosting the volume. |

policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

| Name | Type | Description |
|---|---|---|
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed. |

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

last_error

Error information for the last failed file move on the constituent.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code of the last file move error on the constituent. |

| Name | Type | Description |
| --- | --- | --- |
| destination | integer | DSID of the destination constituent of the last file move error on the constituent. |
| file_id | integer | File ID of the last file move error on the constituent. |
| time | string | Time of the last file move error on the constituent. |

movement

Properties on this constituent related to file movement.

| Name | Type | Description |
| --- | --- | --- |
| file_moves_started | integer | Number of file moves started on this constituent. |
| last_error | last_error | Error information for the last failed file move on the constituent. |
| most_recent_start_time | string | Start time of the most recent file move on the constituent. |

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files. |

| Name | Type | Description |
| --- | --- | --- |
| footprint_invalid | integer | Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints. |
| in_snapshot | integer | Number of blocks skipped by the scanner on this constituent because of files in snapshots. |
| incompatible | integer | Number of blocks skipped by the scanner on this constituent because of incompatible files. |
| metadata | integer | Number of blocks skipped by the scanner on this constituent because of metadata files. |
| on_demand_destination | integer | Number of blocks skipped by the scanner on this constituent because of on demand destination files. |
| other | integer | Number of blocks skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of blocks skipped by the scanner on this constituent because of remote caches. |
| too_large | integer | Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size. |
| too_small | integer | Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of blocks skipped by the scanner on this constituent because of files fenced for write operations. |

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file. |
| footprint_invalid | integer | Number of files skipped by the scanner on this constituent because their space footprints are invalid. |
| in_snapshot | integer | Number of files skipped by the scanner on this constituent because they are trapped in snapshots. |
| incompatible | integer | Number of files skipped by the scanner on this constituent because they are incompatible. |
| metadata | integer | Number of files skipped by the scanner on this constituent because they metadata files. |
| on_demand_destination | integer | Number of files skipped by the scanner on this constituent because they are on demand destinations. |
| other | integer | Number of files skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of files skipped by the scanner on this constituent because they are remote caches. |

| Name | Type | Description |
| --- | --- | --- |
| too_large | integer | Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size. |
| too_small | integer | Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of files skipped by the scanner on this constituent because they are fenced for write operations. |

scanner

Properties related to determining which files to move and where to move them to.

| Name | Type | Description |
| --- | --- | --- |
| blocks_scanned | integer | Number of blocks scanned on this constituent. |
| blocks_skipped | blocks_skipped | Number of blocks skipped by the scanner on this constituent due to various reasons. |
| files_scanned | integer | Number of files scanned on this constituent. |
| files_skipped | files_skipped | Number of files skipped by the scanner on this constituent due to various reasons. |

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | Properties on this constituent related to file movement. |

| Name | Type | Description |
| --- | --- | --- |
| scanner | scanner | Properties related to determining which files to move and where to move them to. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| data_moved | integer | The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent. |
| engine | engine | Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions. |

| Name | Type | Description |
| --- | --- | --- |
| imbalance_percent | integer | Represents the percentage the volume is out of balance. |
| imbalance_size | integer | Represents how much the volume is out of balance, in bytes. |
| max_constituent_imbalance_percent | integer | Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active. |
| notices | array[error] | Capacity rebalancing notice messages. |
| runtime | string | Duration the capacity rebalancing operation has been running. |
| stop_time | string | Time when the capacity rebalancing operation stopped. |
| target_used | integer | Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents. |
| used_for_imbalance | integer | Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size. |

retention

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num> |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |

| Name | Type | Description |
| --- | --- | --- |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | retention | |
| type | string | The SnapLock type of the volume.<br> compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files.<br> enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume.<br> non_snaplock ‐ Indicates the volume is non-snaplock. |
| unspecified_retention_file_count | integer | Indicates the number of files with an unspecified retention time in the volume. |

destinations

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination.<br>• readOnly: 1<br>• Introduced in: 9.9<br>• x-nullable: true |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
| --- | --- | --- |
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |
| used_by_snapshots | integer | Size that is logically used across all snapshots in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

autodelete

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete | autodelete | |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead. |
| reserve_available | integer | Size available for snapshots within the snapshot reserve, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for snapshot usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| auto_adaptive_compression_footprint_data_reduction | integer | Savings achieved due to Auto Adaptive Compression, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| capacity_tier_footprint_data_reduction | integer | Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes. |
| compaction_footprint_data_reduction | integer | Savings achieved due to Data Compaction, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| effective_total_footprint | integer | Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| is_used_stale | boolean | Specifies if the virtual space used is stale. |
| large_size_enabled | boolean | Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |

| Name | Type | Description |
| --- | --- | --- |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. Physical used is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for snapshots from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for snapshot creation, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

access

Raw count and latency data for access operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

dir

Raw count and latency data for directory-create operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

file

Raw count and latency data for file-create operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

symlink

Raw count and latency data for symlink-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

create

Raw count and latency data for create operations.

getattr

Raw count and latency data for getattr operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

link

Raw count and latency data for link operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lock

Raw count and latency data for lock operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lookup

Raw count and latency data for lookup operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

open

Raw count and latency data for open operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

readdir

Raw count and latency data for readdir operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

readlink

Raw count and latency data for readlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

rename

Raw count and latency data for rename operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

setattr

Raw count and latency data for setattr operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

unlink

Raw count and latency data for unlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |

| Name | Type | Description |
| --- | --- | --- |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| cifs_ops_raw | cifs_ops_raw | Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| nfs_ops_raw | nfs_ops_raw | Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM containing the volume. Required on POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tiering

| Name | Type | Description |
| --- | --- | --- |
| min_cooling_days | integer | This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy. |
| object_tags | array[string] | This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value". |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br> all ‐ This policy allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br> auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store<br> none ‐ Volume blocks will not be tiered to the cloud store.<br> snapshot_only ‐ This policy allows tiering of only the volume snapshots not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days. |

| Name | Type | Description |
| --- | --- | --- |
| supported | boolean | This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup volume without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |

| Name | Type | Description |
| --- | --- | --- |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| constituents_per_aggregate | integer | Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |

| Name | Type | Description |
| --- | --- | --- |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |

| Name | Type | Description |
| --- | --- | --- |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| optimize_aggregates | boolean | Specifies whether to create the constituents of the FlexGroup volume on the aggregates specified in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is 'true', the system optimizes the ordering of the aggregates specified. If this value is false, the order of the aggregates is unchanged. The default value is 'false'. |
| qos | qos | QoS information |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |

| Name | Type | Description |
| --- | --- | --- |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |

| Name | Type | Description |
| --- | --- | --- |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | Type of the volume.<br> rw ‐ read-write volume.<br> dp ‐ data-protection volume.<br> ls ‐ load-sharing <code>dp</code> volume. Valid in GET. |

| Name | Type | Description |
| --- | --- | --- |
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a volume

DELETE `/storage/volumes/{uuid}`

**Introduced In:** 9.6

Deletes a volume. If the UUID belongs to a volume, all of its blocks are freed and returned to its containing aggregate. If a volume is online, it is offlined before deletion. If a volume is mounted, unmount the volume by specifying the nas.path as empty before deleting it using the DELETE operation.

#### Optional parameters:

- `force` - Bypasses the recovery-queue and completely removes the volume from the aggregate making it non-recoverable. By default, this flag is set to "false".

#### Related ONTAP commands

- `volume delete`
- `volume clone delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Unique identifier of the volume. |
| force | boolean | query | False | Set the force flag to "true" to bypass the recovery queue, making the deleted volume non-recoverable.<br>• Introduced in: 9.12<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 460770 | [Job "jobid"] Job failed. |
| 524615 | Failed to delete volume because it has one or more clones. |
| 917531 | Cannot delete the root volume because the SVM contains other volumes. |
| 918667 | Volume "name" in SVM "svm.name" cannot be taken offline because it is configured to hold audit logs. |
| 918693 | Volume "name" in SVM "svm.name" cannot be taken offline because it is configured to hold audit logs. |
| 918701 | The specified operation on the volume endpoint is not supported on this platform. |
| 13107349 | Operation is only supported on flexible volumes and FlexGroup volumes. |

| Error Code | Description |
| --- | --- |
| 65537463 | Volume encryption keys (VEK) cannot be created or deleted for data Vserver \"{0}\". External key management has been configured for data Vserver \"{0}\" but ONTAP is not able to encrypt or decrypt with the key manager. Resolve the external key manager key issues at the key manager's portal before creating any new encrypted volumes. |
| 65537529 | Encrypted volumes cannot be created or deleted for Vserver \"{0}\" as a rekey operation for the vserver is in progress. Try creating the encrypted volume again after some time. If the problem persists, run the rekey operation again after some time. |
| 65537600 | Encrypted volumes cannot be created or deleted for Vserver \"{0}\" while the enabled keystore configuration is being switched. If a previous attempt to switch the keystore configuration failed, or was interrupted, the system will continue to prevent encrypted volume creation for Vserver \"{0}\". Use the \"security key-manager keystore enable\" command to re-run and complete the operation. |
| 65539430 | Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation. |
| 65539431 | Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being disabled. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve a volume**

GET `/storage/volumes/{uuid}`

**Introduced In:** 9.6

Retrieves a volume. The GET API can be used to retrieve the quota state for a FlexVol or a FlexGroup volume.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `is_svm_root`
- `analytics.*`
- `anti_ransomware.*`
- `application.*`
- `encryption.*`
- `convert_unicode`
- `clone.parent_snapshot.name`
- `clone.parent_snapshot.uuid`
- `clone.parent_svm.name`
- `clone.parent_svm.uuid`
- `clone.parent_volume.name`
- `clone.parent_volume.uuid`
- `clone.split_complete_percent`
- `clone.split_estimate`
- `clone.split_initiated`
- `efficiency.*`
- `error_state.*`
- `files.*`
- `max_dir_size`
- `nas.export_policy.id`
- `nas.gid`
- `nas.path`
- `nas.security_style`
- `nas.uid`
- `nas.unix_permissions`
- `nas.junction_parent.name`
- `nas.junction_parent.uuid`
- `snaplock.*`

- restore_to.*
- quota.*
- qos.*
- flexcache_endpoint_type
- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint
- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.dedupe_metafiles_footprint
- space.dedupe_metafiles_temporary_footprint
- space.delayed_free_footprint
- space.file_operation_metadata
- space.snapmirror_destination_footprint
- space.volume_guarantee_footprint
- space.cross_volume_dedupe_metafiles_footprint
- space.cross_volume_dedupe_metafiles_temporary_footprint
- space.auto_adaptive_compression_footprint_data_reduction
- space.capacity_tier_footprint_data_reduction
- space.compaction_footprint_data_reduction
- space.effective_total_footprint
- space.snapshot_reserve_unusable
- space.snapshot_spill
- space.user_data
- space.logical_space.*
- space.snapshot.*
- space.used_by_afs
- space.afs_total
- space.available_percent
- space.full_threshold_percent
- space.nearly_full_threshold_percent

- space.overwrite_reserve
- space.overwrite_reserve_used
- space.size_available_for_snapshots
- space.percent_used
- space.fractional_reserve
- space.block_storage_inactive_user_data_percent
- space.physical_used
- space.physical_used_percent
- space.expected_available
- space.filesystem_size
- space.filesystem_size_fixed
- guarantee.*
- autosize.*
- movement.*
- statistics.*
- asynchronous_directory_delete.*
- rebalancing.*
- metric.*
- cloud_write_enabled

#### Related ONTAP commands

- volume show
- volume clone show
- volume efficiency show
- volume encryption show
- volume flexcache show
- volume flexgroup show
- volume move show
- volume quota show
- volume show-space
- volume snaplock show
- volume rebalance show
- security anti-ransomware volume show
- security anti-ransomware volume attack generate-report

- `security anti-ransomware volume space show`
- `volume file async-delete client show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Unique identifier of the volume. |
| is_constituent | boolean | query | False | When set to false, only FlexVol and FlexGroup volumes are returned. When set to true, only FlexGroup constituent volumes are returned. Default for GET calls is false.<br>• Introduced in: 9.10<br>• Default value: |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |

| Name | Type | Description |
| --- | --- | --- |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier.<br>This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |

| Name | Type | Description |
| --- | --- | --- |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
| --- | --- | --- |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| metric | metric | Performance numbers, such as IOPS, latency and throughput. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| statistics | statistics | These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | Type of the volume.<br> rw ‐ read-write volume.<br> dp ‐ data-protection volume.<br> ls ‐ load-sharing <code>dp</code> volume. Valid in GET. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "notices": [
      {
        "code": "124518424",
        "message": "The top metrics report contains partial data for
read operations because NFSv4 reads using Multi-Processor I/O (MPIO)
are not tracked."
      }
    ],
    "state": "string",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
```

```
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
```

```
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "bytes_used": 15436648448,
    "file_count": 21134,
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "report_time": "2024-11-06 13:57:15 -0500",
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
    },
    "state": "string",
    "subdir_count": 35,
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_detected_by": "string",
    "attack_detection_parameters": {
```

```
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "_links": {
          "suspects": {
            "href": "/api/resourcelink"
          }
        },
        "time": "2021-06-01 20:36:41 +0530"
      }
    ],
    "block_device_detection_start_time": "string",
    "block_device_detection_state": "string",
    "clear_suspect": {
      "phase": "string",
      "start_time": "string"
    },
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 23:16:20 +0530"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
```

```
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
      ],
      "historical_statistics": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "surge_statistics": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 23:16:20 +0530"
      },
      "surge_usage": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
```

```
            "name": [
              "lockile"
            ]
          }
        ],
        "time": "2021-12-01 23:16:20 +0530"
      },
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "asynchronous_directory_delete": {
    "trash_bin": "string"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "inherited_physical_used": 0,
    "inherited_savings": 0,
    "lun_name": "string",
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
```

```
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "qtree_name": "string",
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": [
    {
      "aggregates": {
        "name": "string",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "movement": {
        "cutover_window": 30,
        "destination_aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "state": "replicating",
        "tiering_policy": "string"
      },
      "name": "string",
      "node": {
        "name": "string",
        "uuid": "string"
```

```
      },
      "space": {
        "available": 0,
        "block_storage_inactive_user_data": 0,
        "capacity_tier_footprint": 0,
        "footprint": 0,
        "local_tier_footprint": 0,
        "logical_space": {
          "available": 0,
          "used_by_afs": 0
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 19:00:00 +0000",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string",
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
```

```
    "policy": {
      "name": "string"
    },
    "progress": "string",
    "ratio": 0,
    "scanner": {
      "state": "string"
    },
    "schedule": "string",
    "space_savings": {
      "compression": 0,
      "compression_percent": 0,
      "dedupe": 0,
      "dedupe_percent": 0,
      "dedupe_sharing": 0,
      "total": 0,
      "total_percent": 0
    },
    "state": "string",
    "storage_efficiency_mode": "string",
    "type": "string",
    "volume_path": "string"
  },
  "encryption": {
    "action": "string",
    "key_create_time": "2022-01-01 19:00:00 +0000",
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "inodefile_capacity": 0,
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "string",
    "cache_retention_priority": "string",
    "caching_policy": "string"
  },
  "flexcache_endpoint_type": "string",
  "flexgroup": {
    "name": "my_flexgroup",
```

```
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "granular_data_mode": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "cloud": {
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "duration": "PT15S",
    "flexcache": {
      "bandwidth_savings": 4096,
      "cache_miss_percent": 20,
      "duration": "PT1D",
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
```

```
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "movement": {
      "cutover_window": 30,
      "destination_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "percent_complete": 0,
      "start_time": "2020-12-07 03:45:12 -0500",
      "state": "replicating",
      "tiering_policy": "string"
    },
    "name": "vol_cs_dept",
    "nas": {
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 100,
        "name": "default"
      },
      "junction_parent": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "vs1_root",
        "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
      },
      "path": "/user/my_volume",
      "security_style": "string",
```

```
    "unix_permissions": 493
  },
  "nodes": [
    {
      "name": "string",
      "uuid": "string"
    }
  ],
  "optimize_aggregates": true,
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "max_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "quota": {
    "state": "string"
  },
  "rebalancing": {
```

```
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "data_moved": 0,
    "engine": {
      "movement": {
        "file_moves_started": 0,
        "last_error": {
          "code": 0,
          "destination": 0,
          "file_id": 0,
          "time": "2018-06-04 19:00:00 +0000"
        },
        "most_recent_start_time": "2018-06-04 19:00:00 +0000"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
```

```
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      }
    },
    "imbalance_percent": 0,
    "imbalance_size": 0,
    "max_constituent_imbalance_percent": 0,
    "max_runtime": "string",
    "notices": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "start_time": "string",
    "state": "rebalancing",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "sizing_method": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 19:00:00 +0000",
    "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
    "is_audit_log": 1,
    "litigation_count": 10,
    "privileged_delete": "enabled",
    "retention": {
      "default": "P30Y",
      "maximum": "P30Y",
      "minimum": "P30Y"
    },
    "type": "enterprise",
    "unspecified_retention_file_count": 10
```

```
  },
  "snapshot_count": 0,
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 0,
    "cross_volume_dedupe_metafiles_footprint": 0,
    "cross_volume_dedupe_metafiles_temporary_footprint": 0,
    "dedupe_metafiles_footprint": 0,
    "dedupe_metafiles_temporary_footprint": 0,
    "delayed_free_footprint": 0,
    "effective_total_footprint": 0,
    "file_operation_metadata": 0,
    "filesystem_size": 0,
    "footprint": 0,
    "local_tier_footprint": 0,
    "logical_space": {
      "available": 0,
      "used": 0,
      "used_by_afs": 0,
      "used_by_snapshots": 0,
      "used_percent": 0
    },
    "max_size": "string",
    "metadata": 0,
    "over_provisioned": 0,
    "overwrite_reserve": 0,
    "overwrite_reserve_used": 0,
    "percent_used": 0,
    "performance_tier_footprint": 0,
    "size_available_for_snapshots": 0,
    "snapmirror_destination_footprint": 0,
    "snapshot": {
```

```
      "autodelete": {
        "commitment": "string",
        "defer_delete": "string",
        "delete_order": "string",
        "prefix": "string",
        "trigger": "string"
      },
      "autodelete_trigger": "string",
      "reserve_available": 0,
      "reserve_size": 0,
      "space_used_percent": 0,
      "used": 0
    },
    "snapshot_reserve_unusable": 0,
    "snapshot_spill": 0,
    "total_footprint": 0,
    "total_metadata": 0,
    "total_metadata_footprint": 0,
    "used": 0,
    "user_data": 0,
    "volume_guarantee_footprint": 0
  },
  "state": "string",
  "statistics": {
    "cifs_ops_raw": {
      "access": {
        "count": 1000,
        "total_time": 200
      },
      "audit": {
        "count": 1000,
        "total_time": 200
      },
      "create": {
        "dir": {
          "count": 1000,
          "total_time": 200
        },
        "file": {
          "count": 1000,
          "total_time": 200
        },
        "other": {
          "count": 1000,
          "total_time": 200
        },
```

```
        "symlink": {
          "count": 1000,
          "total_time": 200
        }
      },
      "getattr": {
        "count": 1000,
        "total_time": 200
      },
      "link": {
        "count": 1000,
        "total_time": 200
      },
      "lock": {
        "count": 1000,
        "total_time": 200
      },
      "lookup": {
        "count": 1000,
        "total_time": 200
      },
      "open": {
        "count": 1000,
        "total_time": 200
      },
      "read": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
```

```
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
```

```
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
          1055,
          1100,
          700,
          500,
          300,
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
```

```
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      },
      "readdir": {
        "count": 1000,
        "total_time": 200
      },
      "readlink": {
        "count": 1000,
        "total_time": 200
      },
      "rename": {
        "count": 1000,
        "total_time": 200
      },
      "setattr": {
        "count": 1000,
        "total_time": 200
      },
      "unlink": {
        "count": 1000,
        "total_time": 200
      },
      "watch": {
        "count": 1000,
        "total_time": 200
      },
      "write": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
```

```
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
```

```
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
          1055,
          1100,
          700,
          500,
          300,
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
```

```
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      }
    },
    "cloud": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "flexcache_raw": {
      "cache_miss_blocks": 10,
      "client_requested_blocks": 500,
      "status": "ok",
      "timestamp": "2017-01-25 11:20:13 +0000"
    },
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
```

```
      "write": 100
    },
    "nfs_ops_raw": {
      "access": {
        "count": 1000,
        "total_time": 200
      },
      "audit": {
        "count": 1000,
        "total_time": 200
      },
      "create": {
        "dir": {
          "count": 1000,
          "total_time": 200
        },
        "file": {
          "count": 1000,
          "total_time": 200
        },
        "other": {
          "count": 1000,
          "total_time": 200
        },
        "symlink": {
          "count": 1000,
          "total_time": 200
        }
      },
      "getattr": {
        "count": 1000,
        "total_time": 200
      },
      "link": {
        "count": 1000,
        "total_time": 200
      },
      "lock": {
        "count": 1000,
        "total_time": 200
      },
      "lookup": {
        "count": 1000,
        "total_time": 200
      },
      "open": {
```

```
        "count": 1000,
        "total_time": 200
      },
      "read": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
```

```
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
```

```
          1055,
          1100,
          700,
          500,
          300,
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      },
      "readdir": {
        "count": 1000,
        "total_time": 200
      },
      "readlink": {
        "count": 1000,
        "total_time": 200
      },
      "rename": {
        "count": 1000,
        "total_time": 200
      },
      "setattr": {
```

```
        "count": 1000,
        "total_time": 200
      },
      "unlink": {
        "count": 1000,
        "total_time": 200
      },
      "watch": {
        "count": 1000,
        "total_time": 200
      },
      "write": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
```

```
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
```

```
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
          1055,
          1100,
          700,
          500,
          300,
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      }
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
```

```
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "status": [
    "string"
  ],
  "style": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string",
    "supported": true
  },
  "type": "string",
  "use_mirrored_aggregates": true,
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

notices

| Name | Type | Description |
|---|---|---|
| code | string | An error code related to activity tracking. |
| message | string | A notice message related to activity tracking. |

unsupported_reason

| Name | Type | Description |
|---|---|---|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|---|---|---|
| notices | array[notices] | List of activity-tracking notices. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.<br>• enum: ["off", "on"]<br>• Introduced in: 9.10<br>• x-nullable: true |
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

initialization

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the analytics file system scan. |

scan_throttle_reason

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[string] | Arguments present in the warning message encountered. |
| code | string | Warning code indicating why scanner throttling is reported. |
| message | string | A message that provides details for scan throttling. |

unsupported_reason

| Name | Type | Description |
| --- | --- | --- |
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk. |
| file_count | integer | Number of descendants. |
| files_scanned | integer | Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| initialization | initialization | |
| report_time | string | Time of data collection. |
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| scan_throttle_reason | scan_throttle_reason | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume.<br>• enum: ["unknown", "initializing", "initialization_paused", "off", "on"]<br>• Introduced in: 9.8<br>• x-nullable: true |
| subdir_count | integer | Number of sub-directories. |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| total_files | integer | Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`. |
| unsupported_reason | unsupported_reason | |

anti_ransomware_volume_attack_detection_parameters

| Name | Type | Description |
| --- | --- | --- |
| based_on_file_create_op_rate | boolean | Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_delete_op_rate | boolean | Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_rename_op_rate | boolean | Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes. |
| based_on_high_entropy_data_rate | boolean | Specifies whether a high entropy data rate should be considered for attack detection. |
| based_on_never_seen_before_file_extension | boolean | Specifies whether file extensions never seen before should be considered for attack detection. |
| block_device_auto_learned_encryption_threshold | integer | Specifies the block device auto learned encryption threshold. |
| file_create_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file create rate up to which it is considered normal behavior. |
| file_delete_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior. |
| file_rename_op_rate_surge_notify_percent | integer | Specifies the percent of surge in the file rename rate up to which it is considered normal behavior. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_surge_notify_percent | integer | Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on. |
| never_seen_before_file_extension_count_notify_threshold | integer | Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior. |
| never_seen_before_file_extension_duration_in_hours | integer | Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior. |
| relaxing_popular_file_extensions | boolean | Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on. |

_links

| Name | Type | Description |
| --- | --- | --- |
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

clear_suspect

Clear suspect status.

| Name | Type | Description |
|---|---|---|
| phase | string | Clear suspect phase. |
| start_time | string | Clear suspect start time. |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

space

| Name | Type | Description |
|---|---|---|
| snapshot_count | integer | Total number of Anti-ransomware backup snapshots. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup snapshots. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

suspect_files

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume. |
| entropy | string | Indicates the entropy level of this file type. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

historical_statistics

Typical usage values of volume workload.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

newly_observed_file_extensions

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Count of newly observed file extensions. |
| name | string | Name of the newly observed file extension. |

surge_statistics

Usage values of the volume's workload during surge.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

anti_ransomware_volume_workload

| Name | Type | Description |
| --- | --- | --- |
| file_extension_types_count | integer | Count of types of file extensions observed in the volume. |
| file_extensions_observed | array[string] | File extensions observed in the volume. |
| historical_statistics | historical_statistics | Typical usage values of volume workload. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| surge_statistics | surge_statistics | Usage values of the volume's workload during surge. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
| --- | --- | --- |
| attack_detected_by | string | This field specifies whether the attack was reported by `File Analysis` or `Encrypted data percentage analysis`. |
| attack_detection_parameters | anti_ransomware_volume_attack_detection_parameters | |
| attack_probability | string | Probability of a ransomware attack.<br>`none` No suspected ransomware activity.<br>`low` Minimally suspected ransomware activity.<br>`moderate` Moderately suspected ransomware activity.<br>`high` Significantly suspected ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |
| block_device_detection_start_time | string | This field specifies the block device evaluation start time. |
| block_device_detection_state | string | This field specifies the block device attack detection status.<br>`evaluation_period` Attack detection is currently in its evaluation phase.<br>`active_unsuitable_workload` Attack detection is active, but the current workload is not suitable for Anti-ransomware protection.<br>`active_suitable_workload` Attack detection is active, and the current workload is appropriate for Anti-ransomware protection. |
| clear_suspect | clear_suspect | Clear suspect status. |
| dry_run_start_time | string | Time when anti-ransomware monitoring `state` is set to dry-run value for starting evaluation mode.This field will no longer be supported in a future release. |
| event_log | event_log | |

| Name | Type | Description |
| --- | --- | --- |
| space | space | |
| state | string | Anti-ransomware state.<br>`disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.<br>`disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.<br>`dry_run` Anti-ransomware monitoring is enabled in the evaluation mode.<br>`enabled` Anti-ransomware monitoring is active on the volume.<br>`paused` Anti-ransomware monitoring is paused on the volume.<br>`enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.<br>`dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation.<br>For POST, the valid Anti-ransomware states are only `disabled`, `enabled` and `dry_run`, whereas for PATCH, `paused` is also valid along with the three valid states for POST. |
| surge_as_normal | boolean | Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| suspect_files | array[suspect_files] | |

| Name | Type | Description |
| --- | --- | --- |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |
| update_baseline_from_surge | boolean | Sets the observed surge value as the new baseline on a volume. |
| workload | anti_ransomware_volume_workload | |

application

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether asynchronous directory delete from the client is enabled on the volume. |
| trash_bin | string | Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used. |

autosize

| Name | Type | Description |
| --- | --- | --- |
| grow_threshold | integer | Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size.. |
| maximum | integer | Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume. |
| minimum | integer | Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume. |
| mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| shrink_threshold | integer | Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

parent_volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

clone

| Name | Type | Description |
| --- | --- | --- |
| has_flexclone | boolean | Specifies whether this volume is a parent of any FlexClone volume. |
| inherited_physical_used | integer | Inherited physical used from the clone's base snapshot. |
| inherited_savings | integer | Inherited savings from the clone's base snapshot. |
| is_flexclone | boolean | Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST. |
| lun_name | string | This optional parameter specifies the name of a LUN that will be non-disruptively migrated to the newly created FlexClone volume. If not specified, no LUNs are migrated. |
| parent_snapshot | snapshot_reference | |
| parent_svm | parent_svm | SVM, applies only to SVM-scoped objects. |
| parent_volume | parent_volume | |

| Name | Type | Description |
| --- | --- | --- |
| qtree_name | string | This optional parameter specifies the name of the qtree containing the LUN that will be non-disruptively migrated to the newly created FlexClone volume. |
| split_complete_percent | integer | Percentage of FlexClone volume blocks split from its parent volume. |
| split_estimate | integer | Space required by the containing-aggregate to split the FlexClone volume. |
| split_initiated | boolean | This field is set when a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. Valid in PATCH. |

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the aggregate hosting the FlexGroup volume constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the FlexGroup volume constituent. |
| uuid | string | List of the node UUIDs hosting the FlexGroup volume constituent. |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |

| Name | Type | Description |
| --- | --- | --- |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |

| Name | Type | Description |
| --- | --- | --- |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| large_size_enabled | boolean | Specifies whether the support for large volumes and large files is enabled on the volume. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
| --- | --- | --- |
| aggregates | aggregates | |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| name | string | FlexGroup volume constituent name. |
| node | node | |
| space | space | |

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms. |

| Name | Type | Description |
| --- | --- | --- |
| inactive_days | integer | Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'. |
| mode | string | Specifies the mode of inactive data compression scanner. Valid for PATCH and GET. |
| operation_state | string | Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active". |
| status | string | Status of last inactive data compression scan on the volume. |
| threshold_inactive_time | string | Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num> |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform. |

scanner

| Name | Type | Description |
| --- | --- | --- |
| compression | boolean | Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms. |
| dedupe | boolean | Start deduplication if scanning old data. Valid for PATCH and GET. |
| scan_old_data | boolean | Indicates whether or not to scan old data. Valid for PATCH and GET. |
| state | string | State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active". |

space_savings

| Name | Type | Description |
| --- | --- | --- |
| compression | integer | Total disk space that is saved by compressing blocks on the referenced file system, in bytes. |
| compression_percent | integer | Percentage of total disk space that is saved by compressing blocks on the referenced file system. |
| dedupe | integer | Total disk space that is saved by deduplication and file cloning, in bytes. |
| dedupe_percent | integer | Percentage of total disk space that is saved by deduplication and file cloning. |
| dedupe_sharing | integer | Total disk space that is shared due to deduplication and file cloning. |
| total | integer | Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| total_percent | integer | Percentage of total disk space saved in the volume due to deduplication, compression and file cloning. |

efficiency

| Name | Type | Description |
| --- | --- | --- |
| application_io_size | string | Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms. |
| auto_state | string | Automatic deduplication schedule volume state.<br> auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically.<br> deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically. |
| compaction | string | The system can be enabled/disabled compaction.<br> inline ‐ Data will be compacted first and written to the volume.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled. |

| Name | Type | Description |
| --- | --- | --- |
| compression | string | The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be compressed first and written to the volume.<br> background ‐ Data will be written to the volume and compressed later.<br> both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled.<br><br>that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume. |
| compression_type | string | Compression type to use by compression. Valid for PATCH and GET. |

| Name | Type | Description |
| --- | --- | --- |
| cross_volume_dedupe | string | The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be cross volume deduped first and written to the volume.<br> background ‐ Data will be written to the volume and cross volume deduped later.<br> both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled. |
| dedupe | string | The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be deduped first and written to the volume.<br> background ‐ Data will be written to the volume and deduped later.<br> both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled. |
| has_savings | boolean | When true, indicates that the volume contains shared(deduplication, file clones) or compressed data. |

| Name | Type | Description |
| --- | --- | --- |
| idcs_scanner | idcs_scanner | Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| logging_enabled | boolean | When true, indicates that space savings for any newly-written data are being logged. |
| op_state | string | Sis status of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| ratio | number | Storage efficiency that does not include the savings provided by snapshots. |
| scanner | scanner | |
| schedule | string | Schedule associated with volume. |
| space_savings | space_savings | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP and Cloud Volumes ONTAP.<br> disabled ‐ All storage efficiency features are disabled.<br> mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others.<br> On FSx for ONTAP and Cloud Volumes ONTAP ‐<br>   enabled ‐ All supported storage efficiency features for the volume are enabled.<br>   custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled.<br> For other platforms ‐<br>   enabled ‐ At least one storage efficiency feature for the volume is enabled.<br><br>• enum: ["disabled", "enabled", "mixed", "custom"]<br>• Introduced in: 9.9<br>• x-nullable: true |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes. |
| type | string | Sis Type of the volume. |
| volume_path | string | Absolute volume path of the volume. |

status

| Name | Type | Description |
| --- | --- | --- |
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| action | string | This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. The following actions are supported for this field:<br> conversion_pause ‐ Pause an encryption conversion operation currently in progress.<br> conversion_resume ‐ Resume a paused encryption conversion operation.<br> rekey_pause ‐ Pause an encryption rekey operation currently in progress.<br> rekey_resume ‐ Resume a paused encryption rekey operation. |
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_create_time | string | Encryption key creation time of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state.<br> encrypted ‐ The volume is completely encrypted.<br> encrypting ‐ Encryption operation is in progress.<br> partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume.<br> rekeying. Encryption of volume with a new key is in progress.<br> unencrypted ‐ The volume is a plain-text one. |
| status | status | |
| type | string | Volume encryption type.<br> none ‐ The volume is a plain-text one.<br> volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption).<br> aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
| --- | --- | --- |
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |

| Name | Type | Description |
| --- | --- | --- |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies.<br> true ‐ File system is inconsistent.<br> false ‐ File system in not inconsistent. |

files

| Name | Type | Description |
| --- | --- | --- |
| inodefile_capacity | integer | Number of inodes that can currently be stored on the volume for user-visible files. This number dynamically increases as more user-visible files are created. |
| maximum | integer | The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH. |
| used | integer | Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online. |

flash_pool

| Name | Type | Description |
| --- | --- | --- |
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
| --- | --- | --- |
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes

hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| capacity_tier_optimized | boolean | Specifies whether this move should perform optimized volume move for the capacity tier, in which data in the capacity tier does not need to be copied over. |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | Export Policy |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ Unix-style security. |

| Name | Type | Description |
| --- | --- | --- |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth dight), first digit is assumed to be zero. |

nodes

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the volume. |
| uuid | string | List of the node UUIDs hosting the volume. |

policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed. |

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

last_error

Error information for the last failed file move on the constituent.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code of the last file move error on the constituent. |

| Name | Type | Description |
| --- | --- | --- |
| destination | integer | DSID of the destination constituent of the last file move error on the constituent. |
| file_id | integer | File ID of the last file move error on the constituent. |
| time | string | Time of the last file move error on the constituent. |

movement

Properties on this constituent related to file movement.

| Name | Type | Description |
| --- | --- | --- |
| file_moves_started | integer | Number of file moves started on this constituent. |
| last_error | last_error | Error information for the last failed file move on the constituent. |
| most_recent_start_time | string | Start time of the most recent file move on the constituent. |

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files. |

| Name | Type | Description |
| --- | --- | --- |
| footprint_invalid | integer | Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints. |
| in_snapshot | integer | Number of blocks skipped by the scanner on this constituent because of files in snapshots. |
| incompatible | integer | Number of blocks skipped by the scanner on this constituent because of incompatible files. |
| metadata | integer | Number of blocks skipped by the scanner on this constituent because of metadata files. |
| on_demand_destination | integer | Number of blocks skipped by the scanner on this constituent because of on demand destination files. |
| other | integer | Number of blocks skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of blocks skipped by the scanner on this constituent because of remote caches. |
| too_large | integer | Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size. |
| too_small | integer | Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of blocks skipped by the scanner on this constituent because of files fenced for write operations. |

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file. |
| footprint_invalid | integer | Number of files skipped by the scanner on this constituent because their space footprints are invalid. |
| in_snapshot | integer | Number of files skipped by the scanner on this constituent because they are trapped in snapshots. |
| incompatible | integer | Number of files skipped by the scanner on this constituent because they are incompatible. |
| metadata | integer | Number of files skipped by the scanner on this constituent because they metadata files. |
| on_demand_destination | integer | Number of files skipped by the scanner on this constituent because they are on demand destinations. |
| other | integer | Number of files skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of files skipped by the scanner on this constituent because they are remote caches. |

| Name | Type | Description |
| --- | --- | --- |
| too_large | integer | Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size. |
| too_small | integer | Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of files skipped by the scanner on this constituent because they are fenced for write operations. |

scanner

Properties related to determining which files to move and where to move them to.

| Name | Type | Description |
| --- | --- | --- |
| blocks_scanned | integer | Number of blocks scanned on this constituent. |
| blocks_skipped | blocks_skipped | Number of blocks skipped by the scanner on this constituent due to various reasons. |
| files_scanned | integer | Number of files scanned on this constituent. |
| files_skipped | files_skipped | Number of files skipped by the scanner on this constituent due to various reasons. |

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | Properties on this constituent related to file movement. |

| Name | Type | Description |
| --- | --- | --- |
| scanner | scanner | Properties related to determining which files to move and where to move them to. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| data_moved | integer | The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent. |
| engine | engine | Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions. |

| Name | Type | Description |
| --- | --- | --- |
| exclude_snapshots | boolean | Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value. |
| imbalance_percent | integer | Represents the percentage the volume is out of balance. |
| imbalance_size | integer | Represents how much the volume is out of balance, in bytes. |
| max_constituent_imbalance_percent | integer | Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active. |
| max_file_moves | integer | Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value. |

| Name | Type | Description |
| --- | --- | --- |
| max_runtime | string | This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num> |
| max_threshold | integer | Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value. |

| Name | Type | Description |
| --- | --- | --- |
| min_file_size | integer | Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance. |
| min_threshold | integer | Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value. |
| notices | array[error] | Capacity rebalancing notice messages. |
| runtime | string | Duration the capacity rebalancing operation has been running. |

| Name | Type | Description |
| --- | --- | --- |
| start_time | string | Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation.<br><br>While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither.<br><br>If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume.<br><br>If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property.<br><br>The following values apply to FlexGroup volumes.<br> not_running ‐ capacity rebalancing is not running on the volume.<br> starting ‐ used in a PATCH operation to start a capacity rebalancing operation.<br> rebalancing ‐ capacity rebalancing is running on the volume.<br> paused ‐ volume capacity rebalancing is paused on the volume.<br> stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation.<br> unknown ‐ the system was unable to determine the rebalancing state for the volume.<br><br>The following values apply to |

| Name | Type | Description |
| --- | --- | --- |
| stop_time | string | Time when the capacity rebalancing operation stopped. |
| target_used | integer | Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents. |
| used_for_imbalance | integer | Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size. |

retention

the system was unable to determine the rebalancing state for the constituent.

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num> |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |

| Name | Type | Description |
| --- | --- | --- |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | retention | |
| type | string | The SnapLock type of the volume.<br> compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files.<br> enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume.<br> non_snaplock ‐ Indicates the volume is non-snaplock. |
| unspecified_retention_file_count | integer | Indicates the number of files with an unspecified retention time in the volume. |

destinations

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |

| Name | Type | Description |
| --- | --- | --- |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination.<br>• readOnly: 1<br>• Introduced in: 9.9<br>• x-nullable: true |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
| --- | --- | --- |
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |

| Name | Type | Description |
| --- | --- | --- |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |
| used_by_snapshots | integer | Size that is logically used across all snapshots in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

autodelete

| Name | Type | Description |
| --- | --- | --- |
| commitment | string | By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try. |
| defer_delete | string | Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created. |

| Name | Type | Description |
| --- | --- | --- |
| delete_order | string | Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first. |
| enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| prefix | string | Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value. |
| target_free_space | integer | Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized. |
| trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete | autodelete | |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead. |

| Name | Type | Description |
| --- | --- | --- |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead. |
| reserve_available | integer | Size available for snapshots within the snapshot reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for snapshot usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| auto_adaptive_compression_footprint_data_reduction | integer | Savings achieved due to Auto Adaptive Compression, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |

| Name | Type | Description |
| --- | --- | --- |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| capacity_tier_footprint_data_reduction | integer | Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes. |
| compaction_footprint_data_reduction | integer | Savings achieved due to Data Compaction, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| effective_total_footprint | integer | Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| is_used_stale | boolean | Specifies if the virtual space used is stale. |
| large_size_enabled | boolean | Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB. |

| Name | Type | Description |
| --- | --- | --- |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. Physical used is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for snapshots from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for snapshot creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

access

Raw count and latency data for access operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

dir

Raw count and latency data for directory-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

file

Raw count and latency data for file-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

symlink

Raw count and latency data for symlink-create operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

create

Raw count and latency data for create operations.

| Name | Type | Description |
|---|---|---|
| dir | dir | Raw count and latency data for directory-create operations. |
| file | file | Raw count and latency data for file-create operations. |
| other | other | Raw count and latency data for create operations on objects other than files, directories and symlinks. |
| symlink | symlink | Raw count and latency data for symlink-create operations. |

getattr

Raw count and latency data for getattr operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

link

Raw count and latency data for link operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lock

Raw count and latency data for lock operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lookup

Raw count and latency data for lookup operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

open

Raw count and latency data for open operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

readdir

Raw count and latency data for readdir operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

readlink

Raw count and latency data for readlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

rename

Raw count and latency data for rename operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

setattr

Raw count and latency data for setattr operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

unlink

Raw count and latency data for unlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

| Name | Type | Description |
| --- | --- | --- |
| access | access | Raw count and latency data for access operations. |
| audit | audit | Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations. |
| create | create | Raw count and latency data for create operations. |
| getattr | getattr | Raw count and latency data for getattr operations. |
| link | link | Raw count and latency data for link operations. |
| lock | lock | Raw count and latency data for lock operations. |
| lookup | lookup | Raw count and latency data for lookup operations. |
| open | open | Raw count and latency data for open operations. |
| read | read | Raw count and latency data for read operations, including histograms categorizing operations by size and latency. |
| readdir | readdir | Raw count and latency data for readdir operations. |
| readlink | readlink | Raw count and latency data for readlink operations. |
| rename | rename | Raw count and latency data for rename operations. |
| setattr | setattr | Raw count and latency data for setattr operations. |
| unlink | unlink | Raw count and latency data for unlink operations. |

| Name | Type | Description |
| --- | --- | --- |
| watch | watch | Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations. |
| write | write | Raw count and latency data for write operations, including histograms categorizing operations by size and latency. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |

| Name | Type | Description |
| --- | --- | --- |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

| Name | Type | Description |
| --- | --- | --- |
| access | access | Raw count and latency data for access operations. |
| audit | audit | Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations. |
| create | create | Raw count and latency data for create operations. |
| getattr | getattr | Raw count and latency data for getattr operations. |

| Name | Type | Description |
| --- | --- | --- |
| link | link | Raw count and latency data for link operations. |
| lock | lock | Raw count and latency data for lock operations. |
| lookup | lookup | Raw count and latency data for lookup operations. |
| open | open | Raw count and latency data for open operations. |
| read | read | Raw count and latency data for read operations, including histograms categorizing operations by size and latency. |
| readdir | readdir | Raw count and latency data for readdir operations. |
| readlink | readlink | Raw count and latency data for readlink operations. |
| rename | rename | Raw count and latency data for rename operations. |
| setattr | setattr | Raw count and latency data for setattr operations. |
| unlink | unlink | Raw count and latency data for unlink operations. |
| watch | watch | Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations. |
| write | write | Raw count and latency data for write operations, including histograms categorizing operations by size and latency. |

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| cifs_ops_raw | cifs_ops_raw | Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| nfs_ops_raw | nfs_ops_raw | Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM containing the volume. Required on POST.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tiering

| Name | Type | Description |
| --- | --- | --- |
| min_cooling_days | integer | This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy. |
| object_tags | array[string] | This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value". |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br> all ‐ This policy allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br> auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store<br> none ‐ Volume blocks will not be tiered to the cloud store.<br> snapshot_only ‐ This policy allows tiering of only the volume snapshots not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days. |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update volume attributes

PATCH `/storage/volumes/{uuid}`

**Introduced In:** 9.6

Updates the attributes of a volume. For movement, use the "validate_only" field on the request to validate but not perform the operation. The PATCH API can be used to enable or disable quotas for a FlexVol or a FlexGroup volume. The PATCH API can also be used to start or stop non-disruptive volume capacity rebalancing for FlexGroup volumes in addition to modifying capacity rebalancing properties. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path.

A PATCH request for volume encryption performs conversion/rekey operations asynchronously. You can retrieve the conversion/rekey progress details by calling a GET request on the corresponding volume endpoint.

#### Optional properties

- `queue_for_encryption` - Queue volumes for encryption when `encryption.enabled=true`. If this option is not provided or is false, conversion of volumes starts immediately. When there are volumes in the queue and less than four encryptions are running, volumes are encrypted in the order in which they are queued.
- <code>encryption.action</code> - You can pause an ongoing rekey/conversion operation or resume a paused rekey/conversion operation using this field. The following actions are supported for this field: ‐ conversion_pause - Pause an encryption conversion operation currently in progress ‐ conversion_resume - Resume a paused encryption conversion operation ‐ rekey_pause - Pause an encryption rekey operation currently in progress ‐ rekey_resume - Resume a paused encryption rekey operation

#### Related ONTAP commands

- `volume unmount`
- `volume mount`
- `volume online`
- `volume offline`
- `volume modify`
- `volume clone modify`
- `volume efficiency modify`
- `volume quota on`
- `volume quota off`
- `volume snaplock modify`
- `volume encryption conversion start`
- `volume encryption conversion pause`
- `volume encryption conversion resume`
- `volume encryption rekey start`
- `volume encryption rekey pause`

- `volume encryption rekey resume`
- `volume rebalance start`
- `volume rebalance stop`
- `volume rebalance modify`
- `security anti-ransomware volume enable`
- `security anti-ransomware volume disable`
- `security anti-ransomware volume dry-run`
- `security anti-ransomware volume pause`
- `security anti-ransomware volume resume`
- `volume file async-delete client disable`
- `volume file async-delete client enable`
- `volume move`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Unique identifier of the volume. |
| restore_to.snapshot.uuid | string | query | False | UUID of the snapshot to restore volume to the point in time the snapshot was taken. |
| restore_to.snapshot.name | string | query | False | Name of the snapshot to restore volume to the point in time the snapshot was taken. |
| restore_to.path | string | query | False | Path to the file which is restored from the snapshot.<br>• Introduced in: 9.10 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| restore_to.restore_path | string | query | False | Specifies the destination location inside the volume where the file is restored.<br><br>• Introduced in: 9.10 |
| restore_to.start_byte | number | query | False | Starting byte offset of the source file, in multiples of 4096.<br><br>• Introduced in: 9.10 |
| restore_to.byte_count | number | query | False | Number of bytes to restore from the source file, in multiples of 4096.<br><br>• Introduced in: 9.10 |
| preserve_lun_ids | boolean | query | False | Specifies whether LUN IDs need to be preserved during a snapshot restore operation.<br><br>• Introduced in: 9.10<br>• Default value: |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| nvfail | string | query | False | When this option is "on", the filer performs additional work at boot time if it finds that there has been any potential data loss due to an NVRAM failure. In such situations, it causes the invalidation of all NFS file handles on all volumes affected by the problem so that client-side users are forced to remount the affected file system (and thus not continue to use potentially incorrect data). It is also possible to specify a set of files per volume that are renamed in such cases. The filer sends error messages to the console whenever such problems are found.<br>• Introduced in: 9.10<br>• enum: ["off", "on"] |
| snapshot_directory_access_enabled | boolean | query | False | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.<br>• Introduced in: 9.10<br>• Default value: 1 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| sizing_method | string | query | False | Represents the method to modify the size of a Flexgroup. The following methods are supported:<br>• use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources<br>• add_new_resources - Increases the size of the FlexGroup by adding new resources. This is limited to two new resources per available aggregate.<br>• Default value: 1<br>• enum: ["use_existing_resources", "add_new_resources"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| scheduled_snapshot_naming_scheme | string | query | False | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency><br>• Introduced in: 9.10<br>• Default value: 1<br>• enum: ["create_time", "ordinal"]<br><br></scheduled_frequency> |
| validate_only | boolean | query | False | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.<br>• Introduced in: 9.13 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| clone.match_parent_storage_tier | boolean | query | False | Specifies whether the FlexClone volume splits the data blocks by matching its parent storage tier. This option is applicable only if the tiering policy and the tiering minimum cooling days of the parent volume and the FlexClone volume are the same.<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| aggressive_readahead_mode | string | query | False | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. This feature is only available on FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.<br>• Introduced in: 9.13<br>• Default value: 1<br>• enum: ["none", "file_prefetch", "sequential_read", "cross_file_sequential_read"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier.<br>This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| constituents_per_aggregate | integer | Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |

| Name | Type | Description |
| --- | --- | --- |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
| --- | --- | --- |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |

| Name | Type | Description |
| --- | --- | --- |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |

| Name | Type | Description |
| --- | --- | --- |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| tiering | tiering | |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

**Example request**


```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "notices": [
      {
        "code": "124518424",
        "message": "The top metrics report contains partial data for
read operations because NFSv4 reads using Multi-Processor I/O (MPIO)
are not tracked."
      }
    ],
    "state": "string",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
```

```
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
```

```
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "bytes_used": 15436648448,
    "file_count": 21134,
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "report_time": "2024-11-06 13:57:15 -0500",
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
    },
    "state": "string",
    "subdir_count": 35,
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_detected_by": "string",
    "attack_detection_parameters": {
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "time": "2021-06-01 20:36:41 +0530"
      }
    ],
    "block_device_detection_start_time": "string",
    "block_device_detection_state": "string",
```

```
    "clear_suspect": {
      "phase": "string",
      "start_time": "string"
    },
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 23:16:20 +0530"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
      ],
      "historical_statistics": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
```

```
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "surge_statistics": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 23:16:20 +0530"
      },
      "surge_usage": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
            "name": [
              "lockile"
            ]
          }
        ],
        "time": "2021-12-01 23:16:20 +0530"
      },
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
    }
```

```
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "asynchronous_directory_delete": {
    "trash_bin": "string"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "inherited_physical_used": 0,
    "inherited_savings": 0,
    "is_flexclone": true,
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": [
    {
      "aggregates": {
        "name": "string",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "movement": {
        "cutover_window": 30,
        "destination_aggregate": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "state": "replicating",
        "tiering_policy": "string"
      },
      "name": "string",
      "node": {
        "name": "string",
        "uuid": "string"
      },
      "space": {
```

```
        "available": 0,
        "block_storage_inactive_user_data": 0,
        "capacity_tier_footprint": 0,
        "footprint": 0,
        "local_tier_footprint": 0,
        "logical_space": {
          "available": 0,
          "used_by_afs": 0
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 19:00:00 +0000",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string",
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
    "policy": {
      "name": "string"
```

```
      },
      "progress": "string",
      "ratio": 0,
      "scanner": {
        "state": "string"
      },
      "schedule": "string",
      "space_savings": {
        "compression": 0,
        "compression_percent": 0,
        "dedupe": 0,
        "dedupe_percent": 0,
        "dedupe_sharing": 0,
        "total": 0,
        "total_percent": 0
      },
      "state": "string",
      "storage_efficiency_mode": "string",
      "type": "string",
      "volume_path": "string"
    },
    "encryption": {
      "action": "string",
      "key_create_time": "2022-01-01 19:00:00 +0000",
      "key_id": "string",
      "state": "string",
      "status": {
        "code": "string",
        "message": "string"
      },
      "type": "string"
    },
    "files": {
      "inodefile_capacity": 0,
      "used": 0
    },
    "flash_pool": {
      "cache_eligibility": "string",
      "cache_retention_priority": "string",
      "caching_policy": "string"
    },
    "flexcache_endpoint_type": "string",
    "flexgroup": {
      "name": "my_flexgroup",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
```

```
  "granular_data_mode": "string",
  "guarantee": {
    "type": "string"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07 03:45:12 -0500",
    "state": "replicating",
    "tiering_policy": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 493
  },
  "nodes": [
    {
      "name": "string",
      "uuid": "string"
    }
  ],
  "optimize_aggregates": true,
  "qos": {
    "policy": {
      "max_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
```

```
      ],
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput": [
        "900KB/s",
        "500MB/s",
        "120GB/s",
        "5000IOPS",
        "5000IOPS,500KB/s",
        "2500IOPS,100MB/s",
        "1000IOPS,25MB/s"
      ],
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "quota": {
    "state": "string"
  },
  "rebalancing": {
    "data_moved": 0,
    "engine": {
      "movement": {
        "file_moves_started": 0,
        "last_error": {
          "code": 0,
          "destination": 0,
          "file_id": 0,
          "time": "2018-06-04 19:00:00 +0000"
        },
        "most_recent_start_time": "2018-06-04 19:00:00 +0000"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
```

```
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      }
    },
    "imbalance_percent": 0,
    "imbalance_size": 0,
    "max_constituent_imbalance_percent": 0,
    "max_runtime": "string",
    "notices": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "start_time": "string",
    "state": "rebalancing",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
```

```
  "scheduled_snapshot_naming_scheme": "string",
  "sizing_method": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 19:00:00 +0000",
    "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
    "is_audit_log": 1,
    "litigation_count": 10,
    "privileged_delete": "enabled",
    "retention": {
      "default": "P30Y",
      "maximum": "P30Y",
      "minimum": "P30Y"
    },
    "unspecified_retention_file_count": 10
  },
  "snapshot_count": 0,
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 0,
    "cross_volume_dedupe_metafiles_footprint": 0,
    "cross_volume_dedupe_metafiles_temporary_footprint": 0,
    "dedupe_metafiles_footprint": 0,
    "dedupe_metafiles_temporary_footprint": 0,
    "delayed_free_footprint": 0,
    "effective_total_footprint": 0,
    "file_operation_metadata": 0,
    "filesystem_size": 0,
    "footprint": 0,
    "local_tier_footprint": 0,
    "logical_space": {
      "available": 0,
      "used": 0,
      "used_by_afs": 0,
      "used_by_snapshots": 0,
      "used_percent": 0
```

```
      },
      "max_size": "string",
      "metadata": 0,
      "over_provisioned": 0,
      "overwrite_reserve": 0,
      "overwrite_reserve_used": 0,
      "percent_used": 0,
      "performance_tier_footprint": 0,
      "size_available_for_snapshots": 0,
      "snapmirror_destination_footprint": 0,
      "snapshot": {
        "autodelete": {
          "commitment": "string",
          "defer_delete": "string",
          "delete_order": "string",
          "prefix": "string",
          "trigger": "string"
        },
        "autodelete_trigger": "string",
        "reserve_available": 0,
        "reserve_size": 0,
        "space_used_percent": 0,
        "used": 0
      },
      "snapshot_reserve_unusable": 0,
      "snapshot_spill": 0,
      "total_footprint": 0,
      "total_metadata": 0,
      "total_metadata_footprint": 0,
      "used": 0,
      "user_data": 0,
      "volume_guarantee_footprint": 0
    },
    "state": "string",
    "status": [
      "string"
    ],
    "style": "string",
    "tiering": {
      "object_tags": [
        "string"
      ],
      "policy": "string",
      "supported": true
    },
    "use_mirrored_aggregates": true,
```

```
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262256 | Tags cannot be modified for the volume because of the specified reason. |
| 524511 | Failed to restore snapshot "restore_to.snapshot.name". |
| 524849 | Attribute cannot be modified for the target volume because of the specified reason. |
| 787141 | The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates. |
| 787144 | The volume is on an aggregate that is not part of a FabricPool. |

| Error Code | Description |
| --- | --- |
| 917505 | SVM not found. |
| 917534 | The value specified for field "-size" is too small. Update the field "-size" with the minimum size allowed and retry. |
| 917597 | Provisioning a volume on a root aggregate is not supported. |
| 917617 | Cannot modify snapshot policy on volumes of type DC, DP, LS or TMP. |
| 917628 | Modification of the specified fields is not allowed for the target volume type. |
| 917829 | Volume autosize grow threshold must be larger than autosize shrink threshold. |
| 917831 | Volume minimum autosize must be smaller than the maximum autosize. |
| 917888 | Invalid character in volume name. It can have '_' and alphanumeric characters.; |
| 918183 | The specified tiering policy is not supported for the current system configuration. |
| 918193 | Cannot modify tiering min cooling days when vol move is in progress. |
| 918194 | Tiering min cooling days not supported for SVMDR. |
| 918195 | Tiering min cooling days not supported for non data volumes. |
| 918196 | Tiering min cooling days not allowed for the provided tiering policy. |
| 918248 | Specifying a value is not valid for initiating volume FlexClone split operation. |
| 918251 | Specifying a value is not valid for a snapshot restore operation. |
| 918252 | specified "nas.path" is invalid. |
| 918265 | Volume is on the same aggregate. |
| 918266 | "movement.destination_aggregate" and "movement.state" are mutually exclusive, unless the state is "cutover-wait". |
| 918267 | The specified "movement.destination_aggregate" does not exist. |
| 918291 | Invalid volume cloud retrieval policy for the provided tiering policy. |
| 918292 | cloud retrieval policy not supported for non data volume. |

| Error Code | Description |
| --- | --- |
| 918293 | Cannot modify cloud retrieval policy when vol move is in progress. |
| 918334 | Cache cloud retrieval policy requires an effective cluster version of 9.14 or later. |
| 918521 | The volume maximum autosize must be smaller than or equal to the maximum volume size. |
| 918532 | The FlexClone match-parent-storage-tier option requires an effective cluster version of 9.9.1 or later. |
| 918533 | The FlexClone match-parent-storage-tier option not applicable for FlexClone volumes hosted on non-FabricPool storage. |
| 918534 | The tiering policy values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true. |
| 918535 | The tiering minimum cooling day values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true. |
| 918537 | Could not get the FlexClone volume tiering policy or its parent volume tiering policy. Wait a minute and try again. |
| 918538 | The match-parent-storage-tier option is not supported for clone creation. |
| 918577 | Only FlexVol to FlexGroup conversions are currently supported. |
| 918583 | Disabling large volume size is not supported because the volume contains files larger than 16TB or the volume size is larger than 100TB. |
| 918599 | Cloud write requires an effective cluster version of 9.13.1 or later. |
| 918600 | Cloud write is only supported for volumes in a FabricPool on FSx for ONTAP or Cloud Volumes ONTAP. |
| 918601 | Only the \"all\" tiering policy is allowed on a volume with cloud write. |
| 918602 | Cannot move volume when cloud write is enabled. |
| 918603 | Cloud write is only supported on volumes with Read-Write access. |
| 918604 | Cloud write is not supported because the volume is the destination or source endpoint of one or more SnapMirror relationships. |
| 918605 | Cloud write cannot be enabled on a clone volume. |

| Error Code | Description |
| --- | --- |
| 918606 | Cloud write cannot be enabled on a volume containing LUNs. |
| 918607 | Cloud write is not supported because this volume belongs to an SVM that is the source of a DR relationship. |
| 918624 | Tiering storage class requires an effective cluster version of ONTAP 9.13.1 or later. |
| 918625 | Tiering storage class is only supported for volumes in a FabricPool on FSx for ONTAP or Cloud Volumes ONTAP for AWS. |
| 918626 | Tiering storage class is not supported for volumes in a FlexGroup. |
| 918627 | The specified value for "tiering.storage_class" is not supported. |
| 918636 | Volume "name" in SVM "svm.name" is part of a Consistency Group. Snapshot locking cannot be enabled on a volume that is part of a Consistency Group. |
| 918639 | One or more of the fields "-max-size","-max-autosize","-size" are being set or modified. Set the values so that "-max-size" is always greater than or equal to "-max-autosize" and "-size". |
| 918641 | The value passed for "-max-size" is invalid. Valid values for the current platform are specified. |
| 918647 | One or more of the fields "-max-autosize","-size" are being set or modified. Set the values so that "-max-autosize" and "-size" are always less than or equal to the specified value. |
| 918650 | The asynchronous directory delete trash bin name cannot be specified unless asynchronous directory delete from the client is being enabled. |
| 918652 | "error" is an invalid value for field "-state". Valid values are "online", "offline" and "restricted". |
| 918655 | Cloud write is not supported on the current configuration. |
| 918658 | Volume is on a FabricPool-enabled aggregate. CIFS applications might time out when cloud write is enabled on a volume that is on a FabricPool-enabled aggregate. |
| 918659 | Cloud write can only be enabled on FabricPool-enabled aggregates. |

| Error Code | Description |
| --- | --- |
| 918660 | Cloud write enabled volumes can only be moved to FabricPool-enabled aggregates. Disable cloud write to move the volume to a non-FabricPool aggregates. |
| 918703 | Modification of the specified volume properties are not supported on this platform. |
| 918775 | The "constituent_count" parameter must be greater than the current number of constituents in the FlexGroup volume. |
| 918776 | The "constituent_count" and "size" parameters cannot both be specified. |
| 918777 | The "aggregates" parameter cannot contain both storage pods and aggregates. |
| 918786 | The specified value for the "movement.state" parameter is invalid. Valid values are "replicating", "cutover", "aborted", "paused" and "cutover_wait". |
| 1638480 | Failed to promote snapshot "restore_to.snapshot.name" because one or more newer snapshots are currently used as a reference snapshot for data protection operations. |
| 1638590 | Promoting a "-pre-conversion" snapshot is not supported. |
| 2424998 | Unable to determine whether MetroCluster is configured. |
| 2621572 | Modification of the specified field is not permitted on an SVM that is configured as the destination for identity discard SVM DR. |
| 6488103 | Async directory delete from the client is already enabled on volume "name" in SVM "svm.name". |
| 6488104 | Async directory delete from the client is already disabled on volume "name" in SVM "svm.name". |
| 6881612 | Cannot enable inline deduplication on a non AFF Platform or on non HYA aggregate. |
| 6881640 | The "efficiency.state" field is only supported on read-write volumes on FSx for ONTAP and Cloud Volumes ONTAP. |
| 6881651 | File compression is not supported on volumes with **max-size › 300TB**. Modify the volume max-size, then retry. |
| 6881659 | Changing "storage-efficiency-mode" from "efficient" to "default" is not supported. |
| 6881661 | Volume efficiency scanner cannot be started on volume with efficiency policy set as inline-only. |

| Error Code | Description |
| --- | --- |
| 8454210 | Failed to complete the operation because a QoS workload could not be created, modified or deleted. Wait a few seconds, and then try the operation again. |
| 9437885 | The volume is not online. |
| 9437886 | Volume snapshot restore failed because there are newer unexpired locked snapshots in the volume. |
| 13107256 | Operation is only supported on FlexGroup volumes. |
| 13107349 | Operation is only supported on flexible volumes and FlexGroup volumes. |
| 13107371 | Operation is only supported on read-write FlexGroup volumes. |
| 13107404 | When adding new resources to a FlexGroup by specifying "aggregates.name" or "aggregates.uuid", the FlexGroup cannot be resized using "size". These operations must be done separately. |
| 13107415 | Failed to lookup a volume property. |
| 13107431 | Failed to lookup an SVM property. |
| 13107433 | A snapshot is scheduled to be taken within the volume capacity rebalancing runtime. |
| 13107434 | A SnapMirror update is scheduled within the volume capacity rebalancing runtime. |
| 13109187 | When adding new resources to a FlexGroup using "sizing_method", "size" must be specified. Neither "aggregates.name" nor "aggregates.uuid" are allowed to be specified, as the aggregates are selected automatically by the system. |
| 13109198 | Resizing by adding new resources is only supported for FlexGroups. |
| 13109258 | Cannot enable granular data on volume "name" in Vserver "svm.name". This setting can only be enabled on FlexGroups. |
| 13109259 | Granular data cannot be disabled on volume "name" in Vserver "svm.name". This property can only be disabled by restoring a snapshot. |
| 13109260 | Failed to enable granular data on the volume. |
| 13763477 | Cannot modify the SnapLock configuration on volume "name" in SVM "svm.name" because the volume is a part of a consistency group. |

| Error Code | Description |
| --- | --- |
| 65536965 | The key manager on data Vserver \"{0}\" is in the blocked state due to key access errors. Possible reasons for a blocked state include the top-level external key protection key is disabled or not found or the current user does not have access to perform the requested action. Fix the issues with the key manager that manages the top-level external key protection key on data Vserver \"{0}\" to bring the key manager to the active state and then try to bring the volume online again. |
| 65539432 | Cannot online volume \"{1}\" on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation. |
| 65539433 | Cannot online volume \"{1}\" on Vserver \"{0}\" while the keystore is being disabled. |
| 66846758 | Modification of the specified fields is not allowed for FlexCache volumes. |
| 66846772 | Modification of the specified fields is not allowed for origin of FlexCache volumes. |
| 111411201 | File system analytics cannot be enabled on the target volume because of the specified reason. |
| 111411202 | File system analytics cannot be disabled on the target volume because of the specified reason. |
| 111411205 | File system analytics requires an effective cluster version of 9.8 or later. |
| 111411206 | The specified "analytics.state" is invalid. |
| 111411207 | File system analytics cannot be enabled on volumes that contain LUNs. |
| 111411207 | Volume file system analytics is not supported on volumes that contain LUNs. |
| 111411209 | Volume file system analytics is not supported on FlexCache volumes. |
| 111411210 | Volume file system analytics is not supported on audit staging volumes. |
| 111411211 | Volume file system analytics is not supported on object store server volumes. |
| 111411212 | Volume file system analytics is not supported on SnapMirror destination volumes. |
| 111411216 | Enabling or disabling volume file system analytics is not supported on individual FlexGroup constituents. |
| 111411217 | Volume file system analytics is not supported on SnapLock volumes. |

| Error Code | Description |
| --- | --- |
| 111411230 | Volume file system analytics is not supported on volumes that contain NVMe namespaces. |
| 111411241 | Volume file system analytics is not supported for All SAN Array clusters. |
| 111411252 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent`" available space in the volume, then try the operation again. |
| 111411253 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent" available space in all constituents of the FlexGroup, then try the operation again." |
| 111411257 | Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Increase the volume size to "size", then try the operation again." |
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 144180203 | Volume capacity rebalancing is not supported on FlexCache volumes. |
| 144180204 | Volume capacity rebalancing is not supported on object store volumes. |
| 144180207 | Volume capacity rebalancing is not supported on inactive MetroCluster configurations. |

| Error Code | Description |
| --- | --- |
| 144182201 | Volume capacity rebalancing using non-disruptive file move operations and granular data requires an effective cluster version of 9.11.1 or later. |
| 144182203 | The specified value for the "rebalancing.state" parameter is invalid for a PATCH operation. Valid values are "starting" and "stopping". |
| 144182207 | Modifying the volume capacity rebalancing configuration is not supported in the same operation that volume capacity rebalancing is being stopped. |
| 144182211 | The specified value for "-max-file-moves" is not valid. |
| 144182212 | The "-min-file-size" value specified must be larger than 0. |
| 144182213 | The "-min-threshold" value specified must be larger than 0. |
| 144182214 | The "-max-threshold" value specified must be larger than 0. |
| 144182215 | The maximum imbalance threshold value must be larger than the minimum imbalance threshold value. |
| 144182216 | Volume capacity rebalancing is running on the volume. |
| 144182219 | Volume capacity rebalancing for the volume has been started on some constituents. |
| 144182221 | The "-max-runtime" value specified must be 30 minutes or longer. |
| 144182223 | Volume capacity rebalancing is not running on the volume. |
| 144182225 | Internal error in the data component. |
| 144182226 | Failed to load the volume capacity rebalancing configuration for the volume. |
| 144182233 | The specified rebalancing start time must be set to the current time or a later time. |
| 144182235 | The "-min-file-size" value specified must be larger than or equal to 20MB. |
| 144182238 | Volume capacity rebalancing is running on the volume. |
| 196608022 | Failed to start rekey on volume "name" in SVM "svm.name" because of the specified reason. |
| 196608023 | Failed to pause rekey on volume "name" because of the specified reason. |
| 196608024 | Failed to resume rekey on volume "name" because of the specified reason. |

| Error Code | Description |
| --- | --- |
| 196608038 | Failed to start conversion on volume "name" in SVM "svm.name" because of the specified reason. |
| 196608039 | Failed to pause conversion on volume "name" because of the specified reason. |
| 196608041 | Failed to resume conversion on volume "name" because of the specified reason. |
| 196608207 | Rekey operation is already in progress on volume "name". |
| 196608208 | Encryption conversion operation is already in progress on volume "name". |
| 196608209 | Operation is already paused. |
| 203161623 | Disabling anti-ransomware for volume "volume_name" in SVM "vserver_name" is in progress. Wait a few minutes, and try the operation again. |
| 203161629 | Anti-ransomware is not supported on a volume with Snapshot autodelete commitment set to "destroy". |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

notices

| Name | Type | Description |
|---|---|---|
| code | string | An error code related to activity tracking. |
| message | string | A notice message related to activity tracking. |

unsupported_reason

| Name | Type | Description |
|---|---|---|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|---|---|---|
| notices | array[notices] | List of activity-tracking notices. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.<br>• enum: ["off", "on"]<br>• Introduced in: 9.10<br>• x-nullable: true |
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

initialization

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the analytics file system scan. |

scan_throttle_reason

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[string] | Arguments present in the warning message encountered. |
| code | string | Warning code indicating why scanner throttling is reported. |
| message | string | A message that provides details for scan throttling. |

unsupported_reason

| Name | Type | Description |
| --- | --- | --- |
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |

| Name | Type | Description |
| --- | --- | --- |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk. |
| file_count | integer | Number of descendants. |
| files_scanned | integer | Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| initialization | initialization | |
| report_time | string | Time of data collection. |
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`. |
| scan_throttle_reason | scan_throttle_reason | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume.<br>• enum: ["unknown", "initializing", "initialization_paused", "off", "on"]<br>• Introduced in: 9.8<br>• x-nullable: true |
| subdir_count | integer | Number of sub-directories. |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| total_files | integer | Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`. |
| unsupported_reason | unsupported_reason | |

anti_ransomware_volume_attack_detection_parameters

| Name | Type | Description |
|---|---|---|
| based_on_file_create_op_rate | boolean | Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_delete_op_rate | boolean | Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes. |
| based_on_file_rename_op_rate | boolean | Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes. |
| based_on_high_entropy_data_rate | boolean | Specifies whether a high entropy data rate should be considered for attack detection. |
| based_on_never_seen_before_file_extension | boolean | Specifies whether file extensions never seen before should be considered for attack detection. |
| block_device_auto_learned_encryption_threshold | integer | Specifies the block device auto learned encryption threshold. |
| file_create_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file create rate up to which it is considered normal behavior. |
| file_delete_op_rate_surge_notify_percent | integer | Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior. |
| file_rename_op_rate_surge_notify_percent | integer | Specifies the percent of surge in the file rename rate up to which it is considered normal behavior. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_surge_notify_percent | integer | Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on. |
| never_seen_before_file_extension_count_notify_threshold | integer | Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior. |
| never_seen_before_file_extension_duration_in_hours | integer | Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior. |
| relaxing_popular_file_extensions | boolean | Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on. |

anti_ransomware_attack_report

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

clear_suspect

Clear suspect status.

| Name | Type | Description |
|---|---|---|
| phase | string | Clear suspect phase. |
| start_time | string | Clear suspect start time. |

event_log

| Name | Type | Description |
|---|---|---|
| is_enabled_on_new_file_extension_seen | boolean | Specifies whether to send an EMS when a new file extension is discovered. |
| is_enabled_on_snapshot_copy_creation | boolean | Specifies whether to send an EMS when a snapshot is created. |

space

| Name | Type | Description |
|---|---|---|
| snapshot_count | integer | Total number of Anti-ransomware backup snapshots. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup snapshots. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

suspect_files

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume. |
| entropy | string | Indicates the entropy level of this file type. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

historical_statistics

Typical usage values of volume workload.

| Name | Type | Description |
| --- | --- | --- |
| file_create_peak_rate_per_minute | integer | Typical peak rate of file creates per minute in the workload of the volume. |
| file_delete_peak_rate_per_minute | integer | Typical peak rate of file deletes per minute in the workload of the volume. |
| file_rename_peak_rate_per_minute | integer | Typical peak rate of file renames per minute in the workload of the volume. |
| high_entropy_data_write_peak_percent | integer | Typical peak percentage of high entropy data writes in the volume. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Typical peak high entropy data write rate in the volume, in KBs per minute. |

newly_observed_file_extensions

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Count of newly observed file extensions. |
| name | string | Name of the newly observed file extension. |

surge_statistics

Usage values of the volume's workload during surge.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

| Name | Type | Description |
|---|---|---|
| file_create_peak_rate_per_minute | integer | Peak rate of file creates per minute in the workload of the volume during surge. |
| file_delete_peak_rate_per_minute | integer | Peak rate of file deletes per minute in the workload of the volume during surge. |
| file_rename_peak_rate_per_minute | integer | Peak rate of file renames per minute in the workload of the volume during surge. |

| Name | Type | Description |
| --- | --- | --- |
| high_entropy_data_write_peak_percent | integer | Peak percentage of high entropy data writes in the volume during surge. |
| high_entropy_data_write_peak_rate_kb_per_minute | integer | Peak high entropy data write rate in the volume during surge, in KBs per minute. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| time | string | Timestamp at which the first surge in the volume's workload is observed. |

anti_ransomware_volume_workload

| Name | Type | Description |
| --- | --- | --- |
| file_extension_types_count | integer | Count of types of file extensions observed in the volume. |
| file_extensions_observed | array[string] | File extensions observed in the volume. |
| historical_statistics | historical_statistics | Typical usage values of volume workload. |
| newly_observed_file_extensions | array[newly_observed_file_extensions] | New file extensions observed in the volume during surge. |
| surge_statistics | surge_statistics | Usage values of the volume's workload during surge. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
| --- | --- | --- |
| attack_detected_by | string | This field specifies whether the attack was reported by `File Analysis` or `Encrypted data percentage analysis`. |
| attack_detection_parameters | anti_ransomware_volume_attack_detection_parameters | |
| attack_probability | string | Probability of a ransomware attack.<br>`none` No suspected ransomware activity.<br>`low` Minimally suspected ransomware activity.<br>`moderate` Moderately suspected ransomware activity.<br>`high` Significantly suspected ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |
| block_device_detection_start_time | string | This field specifies the block device evaluation start time. |
| block_device_detection_state | string | This field specifies the block device attack detection status.<br>`evaluation_period` Attack detection is currently in its evaluation phase.<br>`active_unsuitable_workload` Attack detection is active, but the current workload is not suitable for Anti-ransomware protection.<br>`active_suitable_workload` Attack detection is active, and the current workload is appropriate for Anti-ransomware protection. |
| clear_suspect | clear_suspect | Clear suspect status. |
| dry_run_start_time | string | Time when anti-ransomware monitoring `state` is set to dry-run value for starting evaluation mode.This field will no longer be supported in a future release. |
| event_log | event_log | |

| Name | Type | Description |
| --- | --- | --- |
| space | space | |
| state | string | Anti-ransomware state.<br>`disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.<br>`disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.<br>`dry_run` Anti-ransomware monitoring is enabled in the evaluation mode.<br>`enabled` Anti-ransomware monitoring is active on the volume.<br>`paused` Anti-ransomware monitoring is paused on the volume.<br>`enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.<br>`dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation.<br>For POST, the valid Anti-ransomware states are only `disabled`, `enabled` and `dry_run`, whereas for PATCH, `paused` is also valid along with the three valid states for POST. |
| surge_as_normal | boolean | Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead. |
| surge_usage | surge_usage | Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead. |
| suspect_files | array[suspect_files] | |

| Name | Type | Description |
| --- | --- | --- |
| typical_usage | typical_usage | Typical usage values of volume workload. This object is no longer supported use historical_statistics instead. |
| update_baseline_from_surge | boolean | Sets the observed surge value as the new baseline on a volume. |
| workload | anti_ransomware_volume_workload | |

application

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies whether asynchronous directory delete from the client is enabled on the volume. |
| trash_bin | string | Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used. |

autosize

| Name | Type | Description |
| --- | --- | --- |
| grow_threshold | integer | Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size.. |
| maximum | integer | Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume. |
| minimum | integer | Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume. |
| mode | string | Autosize mode for the volume.<br> grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value.<br> grow_shrink ‐ Volume grows or shrinks in response to the amount of space used.<br> off ‐ Autosizing of the volume is disabled. |

| Name | Type | Description |
| --- | --- | --- |
| shrink_threshold | integer | Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size. |

snapshot_reference

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

parent_svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

parent_volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

clone

| Name | Type | Description |
| --- | --- | --- |
| has_flexclone | boolean | Specifies whether this volume is a parent of any FlexClone volume. |
| inherited_physical_used | integer | Inherited physical used from the clone's base snapshot. |
| inherited_savings | integer | Inherited savings from the clone's base snapshot. |
| split_complete_percent | integer | Percentage of FlexClone volume blocks split from its parent volume. |
| split_estimate | integer | Space required by the containing-aggregate to split the FlexClone volume. |
| split_initiated | boolean | This field is set when a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. Valid in PATCH. |

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the aggregate hosting the FlexGroup volume constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the FlexGroup volume constituent. |
| uuid | string | List of the node UUIDs hosting the FlexGroup volume constituent. |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |

| Name | Type | Description |
|---|---|---|
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| large_size_enabled | boolean | Specifies whether the support for large volumes and large files is enabled on the volume. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |

| Name | Type | Description |
| --- | --- | --- |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
| --- | --- | --- |
| aggregates | aggregates | |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| name | string | FlexGroup volume constituent name. |
| node | node | |
| space | space | |

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms. |
| inactive_days | integer | Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'. |
| mode | string | Specifies the mode of inactive data compression scanner. Valid for PATCH and GET. |
| operation_state | string | Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active". |
| status | string | Status of last inactive data compression scan on the volume. |
| threshold_inactive_time | string | Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num> |

policy

| Name | Type | Description |
| --- | --- | --- |
| name | string | Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform. |

scanner

| Name | Type | Description |
| --- | --- | --- |
| compression | boolean | Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms. |
| dedupe | boolean | Start deduplication if scanning old data. Valid for PATCH and GET. |
| scan_old_data | boolean | Indicates whether or not to scan old data. Valid for PATCH and GET. |
| state | string | State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active". |

space_savings

| Name | Type | Description |
| --- | --- | --- |
| compression | integer | Total disk space that is saved by compressing blocks on the referenced file system, in bytes. |
| compression_percent | integer | Percentage of total disk space that is saved by compressing blocks on the referenced file system. |
| dedupe | integer | Total disk space that is saved by deduplication and file cloning, in bytes. |
| dedupe_percent | integer | Percentage of total disk space that is saved by deduplication and file cloning. |
| dedupe_sharing | integer | Total disk space that is shared due to deduplication and file cloning. |
| total | integer | Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| total_percent | integer | Percentage of total disk space saved in the volume due to deduplication, compression and file cloning. |

efficiency

| Name | Type | Description |
| --- | --- | --- |
| application_io_size | string | Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms. |
| auto_state | string | Automatic deduplication schedule volume state.<br> auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically.<br> deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically. |
| compaction | string | The system can be enabled/disabled compaction.<br> inline ‐ Data will be compacted first and written to the volume.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled. |

| Name | Type | Description |
| --- | --- | --- |
| compression | string | The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be compressed first and written to the volume.<br> background ‐ Data will be written to the volume and compressed later.<br> both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled.<br><br>that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume. |
| compression_type | string | Compression type to use by compression. Valid for PATCH and GET. |

| Name | Type | Description |
| --- | --- | --- |
| cross_volume_dedupe | string | The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be cross volume deduped first and written to the volume.<br> background ‐ Data will be written to the volume and cross volume deduped later.<br> both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled. |
| dedupe | string | The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms.<br> inline ‐ Data will be deduped first and written to the volume.<br> background ‐ Data will be written to the volume and deduped later.<br> both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.<br> none ‐ None<br> mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled. |
| has_savings | boolean | When true, indicates that the volume contains shared(deduplication, file clones) or compressed data. |

| Name | Type | Description |
| --- | --- | --- |
| idcs_scanner | idcs_scanner | Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| logging_enabled | boolean | When true, indicates that space savings for any newly-written data are being logged. |
| op_state | string | Sis status of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| ratio | number | Storage efficiency that does not include the savings provided by snapshots. |
| scanner | scanner | |
| schedule | string | Schedule associated with volume. |
| space_savings | space_savings | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP and Cloud Volumes ONTAP.<br> disabled ‐ All storage efficiency features are disabled.<br> mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others.<br> On FSx for ONTAP and Cloud Volumes ONTAP ‐<br>   enabled ‐ All supported storage efficiency features for the volume are enabled.<br>   custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled.<br> For other platforms ‐<br>   enabled ‐ At least one storage efficiency feature for the volume is enabled.<br><br>• enum: ["disabled", "enabled", "mixed", "custom"]<br>• Introduced in: 9.9<br>• x-nullable: true |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes. |
| type | string | Sis Type of the volume. |
| volume_path | string | Absolute volume path of the volume. |

status

| Name | Type | Description |
| --- | --- | --- |
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
| --- | --- | --- |
| action | string | This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. The following actions are supported for this field:<br> conversion_pause ‐ Pause an encryption conversion operation currently in progress.<br> conversion_resume ‐ Resume a paused encryption conversion operation.<br> rekey_pause ‐ Pause an encryption rekey operation currently in progress.<br> rekey_resume ‐ Resume a paused encryption rekey operation. |
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_create_time | string | Encryption key creation time of the volume. |

| Name | Type | Description |
| --- | --- | --- |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state.<br> encrypted ‐ The volume is completely encrypted.<br> encrypting ‐ Encryption operation is in progress.<br> partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume.<br> rekeying. Encryption of volume with a new key is in progress.<br> unencrypted ‐ The volume is a plain-text one. |
| status | status | |
| type | string | Volume encryption type.<br> none ‐ The volume is a plain-text one.<br> volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption).<br> aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
| --- | --- | --- |
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |

| Name | Type | Description |
| --- | --- | --- |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies.<br> true ‐ File system is inconsistent.<br> false ‐ File system in not inconsistent. |

files

| Name | Type | Description |
| --- | --- | --- |
| inodefile_capacity | integer | Number of inodes that can currently be stored on the volume for user-visible files. This number dynamically increases as more user-visible files are created. |
| maximum | integer | The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH. |
| used | integer | Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online. |

flash_pool

| Name | Type | Description |
| --- | --- | --- |
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
| --- | --- | --- |
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes

hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
| --- | --- | --- |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation.

| Name | Type | Description |
| --- | --- | --- |
| capacity_tier_optimized | boolean | Specifies whether this move should perform optimized volume move for the capacity tier, in which data in the capacity tier does not need to be copied over. |
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |

| Name | Type | Description |
| --- | --- | --- |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
| --- | --- | --- |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the parent volume. |

nas

| Name | Type | Description |
| --- | --- | --- |
| export_policy | export_policy | Export Policy |
| gid | integer | The UNIX group ID of the volume. Valid in POST or PATCH. |
| junction_parent | junction_parent | |
| path | string | The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline or restricted state removes its junction path. This attribute is reported in GET only when the volume is mounted. |
| security_style | string | Security style associated with the volume. Valid in POST or PATCH.<br> mixed ‐ Mixed-style security<br> ntfs ‐ NTFS/WIndows-style security<br> unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions<br> unix ‐ Unix-style security. |
| uid | integer | The UNIX user ID of the volume. Valid in POST or PATCH. |

| Name | Type | Description |
| --- | --- | --- |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth dight), first digit is assumed to be zero. |

nodes

| Name | Type | Description |
| --- | --- | --- |
| name | string | List of the node names hosting the volume. |
| uuid | string | List of the node UUIDs hosting the volume. |

policy

When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

| Name | Type | Description |
| --- | --- | --- |
| max_throughput | string | Specifies the maximum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either max_throughput_mbps or max_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when max_throughput is set during POST or PATCH. |
| min_throughput | string | Specifies the minimum throughput in Kilobytes per sec, Megabytes per sec or Gigabytes per sec along with or without IOPS. 0 means none.Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when either min_throughput_mbps or min_throughput_iops are set during POST or PATCH. During GET, the returned value is rounded to the largest unit with a value greater than 1. |

| Name | Type | Description |
| --- | --- | --- |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. This cannot be set when min_throughput is set during POST or PATCH. |
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.This cannot be set when min_throughput is set during POST or PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

| Name | Type | Description |
| --- | --- | --- |
| policy | policy | When "min_throughput_iops", "min_throughput_mbps", "min_throughput", "max_throughput_iops", "max_throughput_mbps" or "max_throughput" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed. |

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

last_error

Error information for the last failed file move on the constituent.

| Name | Type | Description |
| --- | --- | --- |
| code | integer | Error code of the last file move error on the constituent. |

| Name | Type | Description |
| --- | --- | --- |
| destination | integer | DSID of the destination constituent of the last file move error on the constituent. |
| file_id | integer | File ID of the last file move error on the constituent. |
| time | string | Time of the last file move error on the constituent. |

movement

Properties on this constituent related to file movement.

| Name | Type | Description |
| --- | --- | --- |
| file_moves_started | integer | Number of file moves started on this constituent. |
| last_error | last_error | Error information for the last failed file move on the constituent. |
| most_recent_start_time | string | Start time of the most recent file move on the constituent. |

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files. |

| Name | Type | Description |
| --- | --- | --- |
| footprint_invalid | integer | Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints. |
| in_snapshot | integer | Number of blocks skipped by the scanner on this constituent because of files in snapshots. |
| incompatible | integer | Number of blocks skipped by the scanner on this constituent because of incompatible files. |
| metadata | integer | Number of blocks skipped by the scanner on this constituent because of metadata files. |
| on_demand_destination | integer | Number of blocks skipped by the scanner on this constituent because of on demand destination files. |
| other | integer | Number of blocks skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of blocks skipped by the scanner on this constituent because of remote caches. |
| too_large | integer | Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size. |
| too_small | integer | Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of blocks skipped by the scanner on this constituent because of files fenced for write operations. |

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

| Name | Type | Description |
| --- | --- | --- |
| efficiency_blocks | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high. |
| efficiency_percent | integer | Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high. |
| fast_truncate | integer | Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file. |
| footprint_invalid | integer | Number of files skipped by the scanner on this constituent because their space footprints are invalid. |
| in_snapshot | integer | Number of files skipped by the scanner on this constituent because they are trapped in snapshots. |
| incompatible | integer | Number of files skipped by the scanner on this constituent because they are incompatible. |
| metadata | integer | Number of files skipped by the scanner on this constituent because they metadata files. |
| on_demand_destination | integer | Number of files skipped by the scanner on this constituent because they are on demand destinations. |
| other | integer | Number of files skipped by the scanner on this constituent for all other reasons. |
| remote_cache | integer | Number of files skipped by the scanner on this constituent because they are remote caches. |

| Name | Type | Description |
| --- | --- | --- |
| too_large | integer | Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size. |
| too_small | integer | Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size. |
| write_fenced | integer | Number of files skipped by the scanner on this constituent because they are fenced for write operations. |

scanner

Properties related to determining which files to move and where to move them to.

| Name | Type | Description |
| --- | --- | --- |
| blocks_scanned | integer | Number of blocks scanned on this constituent. |
| blocks_skipped | blocks_skipped | Number of blocks skipped by the scanner on this constituent due to various reasons. |
| files_scanned | integer | Number of files scanned on this constituent. |
| files_skipped | files_skipped | Number of files skipped by the scanner on this constituent due to various reasons. |

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

| Name | Type | Description |
| --- | --- | --- |
| movement | movement | Properties on this constituent related to file movement. |

| Name | Type | Description |
| --- | --- | --- |
| scanner | scanner | Properties related to determining which files to move and where to move them to. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

| Name | Type | Description |
| --- | --- | --- |
| data_moved | integer | The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent. |
| engine | engine | Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions. |

| Name | Type | Description |
| --- | --- | --- |
| exclude_snapshots | boolean | Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value. |
| imbalance_percent | integer | Represents the percentage the volume is out of balance. |
| imbalance_size | integer | Represents how much the volume is out of balance, in bytes. |
| max_constituent_imbalance_percent | integer | Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active. |
| max_file_moves | integer | Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value. |

| Name | Type | Description |
| --- | --- | --- |
| max_runtime | string | This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num> |
| max_threshold | integer | Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value. |

| Name | Type | Description |
| --- | --- | --- |
| min_file_size | integer | Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance. |
| min_threshold | integer | Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value. |
| notices | array[error] | Capacity rebalancing notice messages. |
| runtime | string | Duration the capacity rebalancing operation has been running. |

| Name | Type | Description |
| --- | --- | --- |
| start_time | string | Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins. |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation.<br><br>While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither.<br><br>If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume.<br><br>If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property.<br><br>The following values apply to FlexGroup volumes.<br> not_running ‐ capacity rebalancing is not running on the volume.<br> starting ‐ used in a PATCH operation to start a capacity rebalancing operation.<br> rebalancing ‐ capacity rebalancing is running on the volume.<br> paused ‐ volume capacity rebalancing is paused on the volume.<br> stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation.<br> unknown ‐ the system was unable to determine the rebalancing state for the volume.<br><br>The following values apply to |

| Name | Type | Description |
| --- | --- | --- |
| stop_time | string | Time when the capacity rebalancing operation stopped. |
| target_used | integer | Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents. |
| used_for_imbalance | integer | Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size. |

retention

the system was unable to determine the rebalancing state for the constituent.

| Name | Type | Description |
| --- | --- | --- |
| default | string | Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| maximum | string | Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

| Name | Type | Description |
| --- | --- | --- |
| minimum | string | Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num> |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| append_mode_enabled | boolean | Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires. |

| Name | Type | Description |
| --- | --- | --- |
| autocommit_period | string | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num> |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |

| Name | Type | Description |
| --- | --- | --- |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | retention | |
| unspecified_retention_file_count | integer | Indicates the number of files with an unspecified retention time in the volume. |

destinations

| Name | Type | Description |
| --- | --- | --- |
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination.<br>• readOnly: 1<br>• Introduced in: 9.9<br>• x-nullable: true |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
| --- | --- | --- |
| destinations | destinations | |

| Name | Type | Description |
| --- | --- | --- |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the snapshot policy.

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

logical_space

| Name | Type | Description |
| --- | --- | --- |
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |
| used_by_snapshots | integer | Size that is logically used across all snapshots in the volume, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

autodelete

| Name | Type | Description |
| --- | --- | --- |
| commitment | string | By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try. |
| defer_delete | string | Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created. |
| delete_order | string | Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first. |
| enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. |
| prefix | string | Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value. |
| target_free_space | integer | Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized. |

| Name | Type | Description |
| --- | --- | --- |
| trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. |

snapshot

| Name | Type | Description |
| --- | --- | --- |
| autodelete | autodelete | |
| autodelete_enabled | boolean | Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead. |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead. |
| reserve_available | integer | Size available for snapshots within the snapshot reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for snapshot usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for snapshot usage, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by snapshots in the volume, in bytes. |

space

| Name | Type | Description |
| --- | --- | --- |
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| auto_adaptive_compression_footprint_data_reduction | integer | Savings achieved due to Auto Adaptive Compression, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| capacity_tier_footprint_data_reduction | integer | Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes. |
| compaction_footprint_data_reduction | integer | Savings achieved due to Data Compaction, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| effective_total_footprint | integer | Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| is_used_stale | boolean | Specifies if the virtual space used is stale. |
| large_size_enabled | boolean | Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| max_size | string | Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume. |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |

| Name | Type | Description |
|---|---|---|
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. Physical used is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for snapshots from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for snapshot creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| total_metadata | integer | Space used by the total metadata in the volume, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| total_metadata_footprint | integer | Space used by the volume metadata footprint in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

access

Raw count and latency data for access operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

dir

Raw count and latency data for directory-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

file

Raw count and latency data for file-create operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

symlink

Raw count and latency data for symlink-create operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

create

Raw count and latency data for create operations.

getattr

Raw count and latency data for getattr operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

link

Raw count and latency data for link operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lock

Raw count and latency data for lock operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

lookup

Raw count and latency data for lookup operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

open

Raw count and latency data for open operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

readdir

Raw count and latency data for readdir operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

readlink

Raw count and latency data for readlink operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

rename

Raw count and latency data for rename operations.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

setattr

Raw count and latency data for setattr operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

unlink

Raw count and latency data for unlink operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

| Name | Type | Description |
|---|---|---|
| count | integer | Number of operations of the given type performed on this volume. |

| Name | Type | Description |
| --- | --- | --- |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• Introduced in: 9.11<br>• x-nullable: true |

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

| Name | Type | Description |
| --- | --- | --- |
| count | integer | Number of operations of the given type performed on this volume. |
| total_time | integer | The raw data component latency in microseconds measured within ONTAP for all operations of the given type.<br>• example: 200<br>• readOnly: 1<br>• x-ntap-advanced: true<br>• Introduced in: 9.11<br>• x-nullable: true |
| volume_protocol_latency_histogram_counts | array[integer] | |
| volume_protocol_latency_histogram_labels | array[string] | Labels for the latency histogram, ranging from <2us to >20s. |
| volume_protocol_size_histogram_counts | array[integer] | |
| volume_protocol_size_histogram_labels | array[string] | Labels for the size histogram, ranging from <4KB to >1024KB. |

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
| --- | --- | --- |
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
| --- | --- | --- |
| cifs_ops_raw | cifs_ops_raw | Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |

| Name | Type | Description |
| --- | --- | --- |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| nfs_ops_raw | nfs_ops_raw | Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |

| Name | Type | Description |
| --- | --- | --- |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM containing the volume. Required on POST.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

tiering

| Name | Type | Description |
| --- | --- | --- |
| min_cooling_days | integer | This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy. |
| object_tags | array[string] | This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value". |

| Name | Type | Description |
| --- | --- | --- |
| policy | string | Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.<br> all ‐ This policy allows tiering of both snapshots and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.<br> auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store<br> none ‐ Volume blocks will not be tiered to the cloud store.<br> snapshot_only ‐ This policy allows tiering of only the volume snapshots not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol volume and "none" for a FlexGroup volume. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _tags | array[string] | Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings. |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |

| Name | Type | Description |
| --- | --- | --- |
| aggregates | array[aggregates] | Aggregate(s) hosting the volume. Optional. |
| aggressive_readahead_mode | string | Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |
| application | application | |
| asynchronous_directory_delete | asynchronous_directory_delete | Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes. |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
| --- | --- | --- |
| cloud_retrieval_policy | string | This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.<br>"on_read" policy retrieves tiered data for all client driven data reads.<br>"never" policy never retrieves tiered data.<br>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved. |
| cloud_write_enabled | boolean | Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools. |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |

| Name | Type | Description |
| --- | --- | --- |
| constituents | array[constituents] | FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid> |
| constituents_per_aggregate | integer | Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |

| Name | Type | Description |
| --- | --- | --- |
| create_time | string | Creation time of the volume. This field is generated when the volume is created. |
| efficiency | efficiency | |
| encryption | encryption | |
| error_state | error_state | |
| files | files | |
| flash_pool | flash_pool | |
| flexcache_endpoint_type | string | FlexCache endpoint type.<br> none ‐ The volume is neither a FlexCache nor origin of any FlexCache.<br> cache ‐ The volume is a FlexCache volume.<br> origin ‐ The volume is origin of a FlexCache volume. |
| flexgroup | flexgroup | |
| granular_data | boolean | State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled.<br>• Introduced in: 9.12<br>• x-nullable: true |

| Name | Type | Description |
| --- | --- | --- |
| granular_data_mode | string | Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled.<br>• enum: ["disabled", "basic", "advanced"]<br>• Introduced in: 9.16<br>• x-nullable: true |
| guarantee | guarantee | |
| has_dir_index_public | boolean | Specifies whether the volume has directories with public index files. |
| has_large_dir | boolean | Specifies whether the volume has directories bigger than 4GB. |
| is_dir_index_transfer_enabled | boolean | When set to true, this field enables support for directory index transfer. |
| is_large_dir_enabled | boolean | When set to true, this field enables the large directory feature on a volume. |

| Name | Type | Description |
| --- | --- | --- |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_s3_arbitrary_part_size_enabled | boolean | Specifies that the volume should allow Amazon S3 multipart uploads with arbitrary part lengths. This is only supported for FlexGroup volumes with `advanced` granular_data. The default value is `false`. When set to `true`, it cannot be reverted to `false`. Clusters with any volumes where this is `true` cannot be reverted to a release that does not support this feature. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| movement | movement | Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. |
| msid | integer | The volume's Master Set ID. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroup volumes, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST. |
| nas | nas | |
| nodes | array[nodes] | List of the nodes hosting the volume. |
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume. |
| rebalancing | rebalancing | Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | Naming Scheme for automatic snapshots:<br>• create_time - Automatic snapshots are saved as per the start of their current date and time.<br>• ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| sizing_method | string | When expanding a FlexGroup volume, this specifies whether to add new constituents, or to resize the current constituents. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of snapshots in the volume. |
| snapshot_directory_access_enabled | boolean | This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. |
| snapshot_locking_enabled | boolean | Specifies whether or not snapshot copy locking is enabled on the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the snapshot policy. |
| space | space | |

| Name | Type | Description |
| --- | --- | --- |
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
| --- | --- | --- |
| style | string | The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol", you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.<br> flexvol ‐ flexible volumes and FlexClone volumes<br> flexgroup ‐ FlexGroup volumes<br> flexgroup_constituent ‐ FlexGroup volume constituents. |
| tiering | tiering | |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

| Name | Type | Description |
|---|---|---|
| validate_only | boolean | Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. |

job_link

| Name | Type | Description |
|---|---|---|
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage files and directories

### Manage files and directories

#### Overview

This API is used to read a file, write to a file, retrieve a list of files and directories, and retrieve or modify certain properties of files and directories. The path field is used to specify the path to the directory or file to be acted on. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

#### File data

Read and write data from/to a named file. To read a file, the Accept request HTTP header must be specified as multipart/form-data, and a value for the `length` query property, which represents the number of bytes to be read, must be specified. The API will fail if the length of data being read/written exceeds 1 MB. This API should only be used on normal files or streams associated with files. The results for other file types, such as LUNs is undefined.

To write to an existing file or create a new file, the user must specify Content-Type as “multipart/form-data”. The “application/json” Content-Type is supported for the files endpoint for file metadata operations. POST, PATCH and GET operations with “application/json” can be used to create a directory, modify metadata fields on a file,

and get directory or file metadata information. POST and PATCH operations created with the Content-Type set to 'application/json' are not supported for use with the *data* field.
The following APIs are used to read or write data to a file:

– GET /api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=0&length=40 -H "Accept: multipart/form-data"

– POST /api/storage/volumes/{volume.uuid}/files/{path} -H "Content-Type: multipart/form-data" --form "file=the data to be written to the new file"

– PATCH /api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=10 -H "Content-Type: multipart/form-data" --form "file=the new data to be written or overwritten to the existing file starting at byte_offset"

#### Listing directories and files

A list of files and directories and their properties can be retrieved for a specified path.

The following APIs are used to view a list of files and directories:

– GET /api/storage/volumes/{volume.uuid}/files

– GET /api/storage/volumes/{volume.uuid}/files/{path}

– GET /api/storage/volumes/{volume.uuid}/files/{path}?fields=*

#### File information

The metadata and detailed information about a single directory or file can be retrieved by setting the `return_metadata` query property to `true`. The information returned includes type, creation_time, modified_time, changed_time, accessed_time, unix_permissions, owner_id, group_id, size, hard_links_count, inode_number, is_empty, bytes_used, inode_generation, is_vm_aligned, is_junction, links, and analytics (if requested).

The following API is used to view the properties of a single file or directory:

– GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true

#### File usage

Custom details about the usage of a file can be retrieved by specifying a value for the `byte_offset` and `length` query properties.

The following API is used to view the bytes used, by a file based on the range defined by `byte_offset` and `length`:

– GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&byte_offset={int}&length={int}

#### Create a directory

The following API is used to create a directory:

– POST /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "type" : "directory", "unix-permissions" : "644"}'

#### Delete an entire directory

A directory can be deleted. The behavior of this call is equivalent to rm -rf.

The following API is used to delete an entire directory:

– DELETE /api/storage/volumes/{volume.uuid}/files/{path}?recurse=true

#### Delete a file or an empty directory

The following API is used to delete a file or an empty directory:

– DELETE /api/storage/volumes/{volume.uuid}/files/{path}

– DELETE /api/storage/volumes/{volume.uuid}/files/{path}?recurse=false

#### File system analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. When file system analytics are enabled on a volume, `analytics.*` fields may be requested, and will be populated in the response records corresponding to directories. The API does not support file system analytics for requests that are made beyond the boundary of the specified `volume.uuid`.

The following APIs are used to obtain analytics information for a directory:

– GET /api/storage/volumes/{volume.uuid}/files/{path}?fields=analytics

– GET /api/storage/volumes/{volume.uuid}/files/{path}?fields=**

#### QoS

QoS policies and settings enforce Service Level Objectives (SLO) on a file. A pre-created QoS policy can be used by specifying the `qos.name` or `qos.uuid` properties.

The following APIs are used to assign a QoS policy to a file:

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "qos_policy.name" : "policy" }'

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "qos_policy.uuid" : "b89bc5dd-94a3-11e8-a7a3-0050568edf84" }'

#### Symlinks

The following APIs are used to create a symlink and read the contents of a symlink:

– POST /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "target" : "directory2/file1" }'

– GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&fields=target

#### Rename a file or a directory

The following API can be used to rename a file or a directory. Note that you need to provide the path relative to the root of the volume in the `path` body parameter.

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "path" : "directory1/directory2" }'

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "path" : "directory1/directory2/file1" }'

#### Examples

### Writing to a new file

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile" -H "Content-Type: multipart/form-data"
--form "file=the data to be written to the new file"
```

### Writing to an existing file

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile?byte_offset=39" -H "Content-Type:
multipart/form-data" --form "file=*here is a little more data"
```

### Reading a file

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile?byte_offset=0&length=100" -H "Accept:
multipart/form-data"

# Response for file data:
--ec51b3541741ade7
Content-Disposition: form-data; name="bytes_read"
Content-Type: text/plain
66
--ec51b3541741ade7
Content-Disposition: form-data; filename="aNewFile"
Content-Type: application/octet-stream
the data to be written to the new file*here is a little more data
--ec51b3541741ade7--
```

### Creating a directory

You can use the POST request to create a directory.

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/dir1"  -H 'accept: application/hal+json' -d '{
"type" : "directory", "unix_permissions" : "644" }'

# The response:
{
"num_records": 1,
"records": [
  {
    "path": "dir1",
    "type": "directory",
    "unix_permissions": 644
  }
]
 }
```

### Creating a stream on a file

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}?overwrite=true

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile?overwrite=true&byte_offset=-
1&stream_name=someStream" -H "Content-Type: multipart/form-data" --form
"file=the data to be written to the new file"
```

### Retrieving the list of files in a directory

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/d1%2Fd2%2Fd3"
```

```
# Response for file records:
{
"records": [
  {
    "path": "d1/d2/d3",
    "name": ".",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "..",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "f1",
    "type": "file",
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2File1?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
```

```
    "name": "d5",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3"
  }
}
}
```

#### Retrieving a list of files based on file type

You can filter the list of files you retrieve based on multiple file types by including a query parameter in the following format type="file\|symlink"

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/d1%2Fd2%2Fd3?type=file&#124;directory"

# Response for file records:
{
"records": [
  {
    "path": "d1/d2/d3",
    "name": ".",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
```

```
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "..",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "f1",
    "type": "file",
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2File1?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "d5",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
      }
```

```
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3"
  }
}
}
```


**Retrieving the properties of a directory or a file**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"

# Response for file properties:
{
"records": [
  {
    "path": "d1/d2/d3/f1",
    "name": "",
    "type": "file",
    "creation_time": "2019-06-12T21:27:28-04:00",
    "modified_time": "2019-06-12T21:27:28-04:00",
    "changed_time": "2019-06-12T21:27:28-04:00",
    "accessed_time": "2019-06-12T21:27:28-04:00",
    "unix_permissions": 644,
    "owner_id": 54738,
    "group_id": 30,
    "size": 200,
    "hard_links_count": 1,
    "inode_number": 1233,
    "bytes_used": 4096,
    "inode_generation": 214488325,
    "is_vm_aligned": false,
    "is_junction": false
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/da8bb06c-823e-11e9-b790-
005056acdcb0/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"
  }
}
}
```

#### Creating a symlink to a relative path

You can use the POST request to create a symlink.

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/symlink1"  -H 'accept: application/hal+json' -d '{
"target" : "d1/f1"}'

# The response:
{
"num_records": 1,
"records": [
  {
    "path": "symlink1",
    "target": "d1/f1"
  }
]
}
```

#### Retrieving the target of a symlink

You can use the GET request to view the target of a symlink.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/symlink1?return_metadata=true&fields=target"

# The response:
{
"records": [
  {
    "path": "symlink1",
    "target": "d1/f1"
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/54c06ce2-5430-11ea-90f9-
005056a73aff/files/symlink1?return_metadata=true&fields=target"
  }
}
}
```

#### Retrieving the usage information for a file

You can use the GET request to retrieve the bytes used in a file with or without specifying the offset.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"

# The response:
{
"records": [
  {
    "path": "d1/d2/d3/f1",
    "type": "file",
    "creation_time": "2019-06-12T21:27:28-04:00",
    "modified_time": "2019-06-12T21:27:28-04:00",
    "changed_time": "2019-06-12T21:27:28-04:00",
    "accessed_time": "2019-06-12T21:27:28-04:00",
    "unix_permissions": 644,
    "owner_id": 54738,
    "group_id": 30,
    "size": 200,
    "hard_links_count": 1,
    "inode_number": 1233,
    "bytes_used": 4096,
    "inode_generation": 214488325,
    "is_vm_aligned": false,
    "is_junction": false
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b139-8d21-11e9-b926-
05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"
  }
}
}
```

**Retrieving all information (including analytics) for a directory**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-
8043-00505682f860/files/d1?return_metadata=true&fields=**"

# Response for all fields of the directory:
{
"records": [
  {
    "svm": {
      "uuid": "58a996a2-f9d5-11e9-8043-00505682f860",
      "_links": {
        "self": {
          "href": "/api/svm/svms/58a996a2-f9d5-11e9-8043-00505682f860"
        }
      }
    },
    "volume": {
        "uuid": "1ef5d1b2-f9d7-11e9-8043-00505682f860",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-
00505682f860"
          }
      }
    },
    "path": "d1",
    "type": "directory",
    "creation_time": "2019-10-28T23:04:13+00:00",
    "modified_time": "2019-10-28T23:10:30+00:00",
    "changed_time": "2019-10-28T23:10:30+00:00",
    "accessed_time": "2019-10-28T23:10:38+00:00",
    "unix_permissions": 755,
    "owner_id": 1002,
    "group_id": 65533,
    "size": 4096,
    "hard_links_count": 5,
    "inode_number": 96,
    "is_empty": false,
    "bytes_used": 4096,
    "inode_generation": 214514951,
    "is_vm_aligned": false,
    "is_junction": false,
    "analytics": {
      "file_count": 668,
      "bytes_used": 209657856,
      "subdir_count": 18,
      "by_modified_time": {
```

```
        "bytes_used": {
          "values": [
            0,
            0,
            0,
            0,
            3112960,
            0,
            14041088,
            20545536,
            0,
            57933824,
            61947904,
            68804608,
            188686336,
            0,
            0,
            0,
            20971520,
            0
          ],
          "percentages": [
            0,
            0,
            0,
            0,
            1.48,
            0,
            6.7,
            9.8,
            0,
            27.63,
            29.55,
            32.82,
            90,
            0,
            0,
            0,
            10,
            0
          ],
          "labels": [
            "2019-W42",
            "2019-W41",
            "2019-W40",
            "2019-W39",
```

```
            "2019-W38",
            "2019-10",
            "2019-09",
            "2019-08",
            "2019-Q4",
            "2019-Q3",
            "2019-Q2",
            "2019-Q1",
            "2019",
            "2018",
            "2017",
            "2016",
            "--2015",
            "unknown"
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
          "values": [
            102760448,
            1867776,
            1245184,
            2179072,
            1556480,
            105873408,
            9027584,
            8093696,
            105873408,
            23969792,
            32382976,
            26460160,
            188686336,
            0,
            0,
            0,
            20971520,
            0
          ],
          "percentages": [
            49.01,
            0.89,
            0.59,
            1.04,
            0.74,
            50.5,
```

```
            4.31,
            3.86,
            50.5,
            11.43,
            15.45,
            12.62,
            90,
            0,
            0,
            0,
            10,
            0
          ],
          "labels": [
            "2019-W42",
            "2019-W41",
            "2019-W40",
            "2019-W39",
            "2019-W38",
            "2019-10",
            "2019-09",
            "2019-08",
            "2019-Q4",
            "2019-Q3",
            "2019-Q2",
            "2019-Q1",
            "2019",
            "2018",
            "2017",
            "2016",
            "--2015",
            "unknown"
          ]
        }
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-
00505682f860/files/d1?return_metadata=true&fields=**"
  }
}
}
```

**Retrieving file system analytics information for a set of histogram buckets**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-
05056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_b
y_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"

# Response with analytics data
{
"records": [
  {
    "path": "d3",
    "name": ".",
    "type": "directory",
    "analytics": {
      "file_count": 44,
      "bytes_used": 244240384,
      "subdir_count": 14,
      "by_modified_time": {
        "bytes_used": {
          "values": [
            57344,
            29720576,
            196141056,
            57344
          ],
          "percentages": [
            0.02,
            12.17,
            80.31,
            0.02
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
          "values": [
            69632,
            244170752,
            0,
            0
          ],
```

```
          "percentages": [
            0.03,
            99.97,
            0,
            0
          ]
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2F%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2F%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d3",
    "name": "..",
    "type": "directory",
    "analytics": {
      "file_count": 515,
      "bytes_used": 3034574848,
      "subdir_count": 23,
      "by_modified_time": {
        "bytes_used": {
          "values": [
            61440,
            1756479488,
            214622208,
            1191936
          ],
          "percentages": [
            0,
            57.88,
            7.07,
            0.04
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
```

```
          "values": [
            282624,
            3034292224,
            0,
            0
          ],
          "percentages": [
            0.01,
            99.99,
            0,
            0
          ]
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d3",
    "name": "d5",
    "type": "directory",
    "analytics": {
      "file_count": 10,
      "bytes_used": 47648768,
      "subdir_count": 4,
      "by_modified_time": {
        "bytes_used": {
          "values": [
            0,
            29638656,
            0,
            0
          ],
          "percentages": [
            0,
            62.20,
            0,
```

```
            0
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
          "values": [
            0,
            47648768,
            0,
            0
          ],
          "percentages": [
            0,
            100,
            0,
            0
          ]
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2Fd5"
      },
      "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3%2Fd5?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"analytics": {
  "by_modified_time": {
    "bytes_used": {
      "labels": [
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2018-Q4"
      ]
    }
  },
  "by_accessed_time": {
```

```
    "bytes_used": {
      "labels": [
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2018-Q4"
      ]
    }
  }
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_
by_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"
  }
}
}
```

**Identifying the largest subdirectories**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-
8043-
00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=
analytics.bytes_used%20desc"

# Response with the largest subdirectories sorted by their usage:
{
"records": [
  {
    "path": "d1",
    "name": "..",
    "type": "directory",
    "analytics": {
      "bytes_used": 56623104
    }
  },
  {
    "path": "d1",
    "name": ".",
    "type": "directory",
```

```
    "analytics": {
      "bytes_used": 35651584
    }
  },
  {
    "path": "d1",
    "name": "biggest",
    "type": "directory",
    "analytics": {
      "bytes_used": 17825792
    }
  },
  {
    "path": "d1",
    "name": "bigger",
    "type": "directory",
    "analytics": {
      "bytes_used": 10485760
    }
  },
  {
    "path": "d1",
    "name": "big",
    "type": "directory",
    "analytics": {
      "bytes_used": 5242880
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-
00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=
analytics.bytes_used%20desc"
  }
}
}
```


#### Assigning a QoS policy to a file

You can use the PATCH request to assign a QoS policy to a file.

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Ffile1" -d '{ "qos_policy": { "name":
"policy" }}'

# The response:
{}
```

#### Retrieving QoS information for a file

You can use the GET request for all fields with return_metadata="true" to retrieve QoS information for the file.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/file?return_metadata=true&fields=**"

# The response:
{
"records": [
  {
    "svm": {
     "uuid": "42ee3002-67dd-11ea-8508-005056a7b8ac"
    },
    "volume": {
      "uuid": "c05eb66a-685f-11ea-8508-005056a7b8ac"
    },
    "path": "file",
    "type": "lun",
    "creation_time": "2020-03-17T10:58:40-04:00",
    "modified_time": "2020-03-24T18:15:40-04:00",
    "changed_time": "2020-03-24T18:15:40-04:00",
    "accessed_time": "2020-03-24T18:15:40-04:00",
    "unix_permissions": 644,
    "owner_id": 0,
    "group_id": 0,
    "size": 1048576,
    "hard_links_count": 2,
    "inode_number": 96,
    "bytes_used": 1056768,
    "inode_generation": 219748425,
    "is_vm_aligned": false,
    "is_junction": false,
    "is_snapshot": false,
    "qos_policy": {
      "name": "pg1",
      "uuid": "00725264-688f-11ea-8f10-005056a7b8ac"
    }
  }
] ,
"num_records": 1
}
```

### Deleting an entire directory

You can use the DELETE request to remove an entire directory recursively.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Fdirectory2?recurse=true"

# The response:
{
"job": {
  "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"
    }
  }
}
}
```

### Deleting an entire directory with specified throttling threshold

You can specify the maximum number of directory delete operations per second when removing an entire directory recursively.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-
05056aca658/files/directory1%2Fdirectory2?recurse=true&throttle_deletion=1
00"

# The response:
{
"job": {
  "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"
    }
  }
}
}
```

#### Deleting an empty directory

You can use the DELETE request to remove an empty directory.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Fdirectory2"

# The response:
{}
```

#### Deleting a file

You can use the DELETE request to remove a file.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Ffile2"

# The response:
{}
```

#### Renaming a file

You can use the PATCH request to rename a file.

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Fdirectory2%2Ffile1" -d '{ "path":
"directory1/file2" }'

# The response:
{}
```

#### File truncating

You can use the PATCH request to change the size of a file.

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/abc.txt" -d '{ "size":100 }'

# The response:
{}
```

#### Renaming a directory

You can use the PATCH request to rename a directory.

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Fdirectory2" -d '{ "path": "d3/d4" }'

# The response:
{}
```

### Delete an existing file or directory

DELETE `/storage/volumes/{volume.uuid}/files/{path}`

**Introduced In:** 9.8

Deletes an existing file or directory. Query-based DELETE operations are not supported.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| path | string | path | True | The relative path of a directory in the volume. The path field requires using "%2E" to represent ".", "%2E%2E" to represent ".." and "%2F" to represent "/" for the path provided. Refer to "relative reference resolution" section in Getting started with the ONTAP REST API for considerations when using "..". |
| recurse | boolean | query | False | Delete an entire directory. The behavior of this call is equivalent to rm -rf.<br>• Default value: |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throttle_deletion | integer | query | False | The maximum number of directory delete operations per second. A valid throttle_deletion number is an integer from 10 to 100000. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Response**

```
Status: 202, Accepted
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 131102 | Read-only file system. |
| 131138 | Directory not empty. |
| 918235 | A volume with the specified UUID was not found. |
| 6488081 | The {field} field is not supported for DELETE operations. |
| 6488110 | A volume delete is not supported on this endpoint. |
| 6684674 | No such file or directory. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve files and directories

GET `/storage/volumes/{volume.uuid}/files/{path}`

**Introduced In:** 9.7

Retrieves a list of files and directories for a given directory or returns only the properties of a single given directory or file of a volume.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query property. See Requesting specific fields to learn more.

- `analytics`
- `qos_policy.name`
- `qos_policy.uuid`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| path | string | path | True | Relative path of a file or directory in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided. |
| byte_offset | integer | query | False | The file offset to start reading from.<br>• Introduced in: 9.8 |
| length | integer | query | False | Length of the range in bytes.<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| return_metadata | boolean | query | False | If true, the request returns metadata for the the directory or file specified in the path.<br>• Introduced in: 9.8<br>• Default value: |
| unix_permissions | integer | query | False | Filter by unix_permissions |
| creation_time | string | query | False | Filter by creation_time |
| target | string | query | False | Filter by target<br>• Introduced in: 9.8 |
| type | string | query | False | Filter by type |
| inode_number | integer | query | False | Filter by inode_number |
| name | string | query | False | Filter by name |
| inode_generation | integer | query | False | Filter by inode_generation |
| constituent.uuid | string | query | False | Filter by constituent.uuid<br>• Introduced in: 9.10 |
| constituent.name | string | query | False | Filter by constituent.name<br>• Introduced in: 9.10 |
| size | integer | query | False | Filter by size |
| is_empty | boolean | query | False | Filter by is_empty |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| analytics.by_modified_time.bytes_used.labels | string | query | False | Filter by analytics.by_modified_time.bytes_used.labels<br><br>• Introduced in: 9.8 |
| analytics.by_modified_time.bytes_used.newest_label | string | query | False | Filter by analytics.by_modified_time.bytes_used.newest_label<br><br>• Introduced in: 9.8 |
| analytics.by_modified_time.bytes_used.oldest_label | string | query | False | Filter by analytics.by_modified_time.bytes_used.oldest_label<br><br>• Introduced in: 9.8 |
| analytics.by_modified_time.bytes_used.values | integer | query | False | Filter by analytics.by_modified_time.bytes_used.values<br><br>• Introduced in: 9.8 |
| analytics.by_modified_time.bytes_used.aged_data_metric | number | query | False | Filter by analytics.by_modified_time.bytes_used.aged_data_metric<br><br>• Introduced in: 9.17 |
| analytics.by_modified_time.bytes_used.percentages | number | query | False | Filter by analytics.by_modified_time.bytes_used.percentages<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| analytics.subdir_count | integer | query | False | Filter by analytics.subdir_count<br><br>• Introduced in: 9.8 |
| analytics.report_time | string | query | False | Filter by analytics.report_time<br><br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.labels | string | query | False | Filter by analytics.by_accessed_time.bytes_used.labels<br><br>• Introduced in: 9.8 |
| analytics.by_accessed_time.bytes_used.newest_label | string | query | False | Filter by analytics.by_accessed_time.bytes_used.newest_label<br><br>• Introduced in: 9.8 |
| analytics.by_accessed_time.bytes_used.oldest_label | string | query | False | Filter by analytics.by_accessed_time.bytes_used.oldest_label<br><br>• Introduced in: 9.8 |
| analytics.by_accessed_time.bytes_used.values | integer | query | False | Filter by analytics.by_accessed_time.bytes_used.values<br><br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| analytics.by_accessed_time.bytes_used.aged_data_metric | number | query | False | Filter by analytics.by_accessed_time.bytes_used.aged_data_metric<br>• Introduced in: 9.17 |
| analytics.by_accessed_time.bytes_used.percentages | number | query | False | Filter by analytics.by_accessed_time.bytes_used.percentages<br>• Introduced in: 9.8 |
| analytics.bytes_used | integer | query | False | Filter by analytics.bytes_used<br>• Introduced in: 9.8 |
| analytics.incomplete_data | boolean | query | False | Filter by analytics.incomplete_data<br>• Introduced in: 9.12 |
| analytics.file_count | integer | query | False | Filter by analytics.file_count<br>• Introduced in: 9.8 |
| is_snapshot | boolean | query | False | Filter by is_snapshot<br>• Introduced in: 9.8 |
| qos_policy.uuid | string | query | False | Filter by qos_policy.uuid<br>• Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| qos_policy.name | string | query | False | Filter by qos_policy.name<br>• Introduced in: 9.8 |
| overwrite_enabled | boolean | query | False | Filter by overwrite_enabled<br>• Introduced in: 9.8 |
| volume.name | string | query | False | Filter by volume.name |
| group_id | integer | query | False | Filter by group_id |
| bytes_used | integer | query | False | Filter by bytes_used |
| owner_id | integer | query | False | Filter by owner_id |
| is_vm_aligned | boolean | query | False | Filter by is_vm_aligned |
| hard_links_count | integer | query | False | Filter by hard_links_count |
| unique_bytes | integer | query | False | Filter by unique_bytes<br>• Introduced in: 9.8 |
| fill_enabled | boolean | query | False | Filter by fill_enabled<br>• Introduced in: 9.8 |
| changed_time | string | query | False | Filter by changed_time |
| accessed_time | string | query | False | Filter by accessed_time |
| modified_time | string | query | False | Filter by modified_time |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| is_junction | boolean | query | False | Filter by is_junction |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| analytics.histogram_by_time_labels | array[string] | query | False | Request that returned analytics_histogram_by_time objects including values associated with the specified labels.<br><br>As described in the object description, the labels may take the following forms:<ul><br>partial-date<br><tt>--</tt>partial-date<br>partial-date <tt>--</tt><br>partial-date<tt>-</tt>-partial-date<br><tt>unknown</tt><br><br></ul>Intervals that the system would not normally return may be specified. In this case, the appropriate values and percentages summarizing all files with a time-based attribute within the indicated period of time are calculated and returned in the response. However, there are some restrictions:<ul><br> Any partial-date specified as the beginning or end of an interval must be tracked by the system. Valid partial-dates may be determined by making an OPTIONS request to |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| analytics | analytics | Additional file system analytics information that is invariant amongst all elements in the collection.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>This analytics object captures properties that are invariant amongst all elements included in the `records` array. The invariant properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large. |
| num_records | integer | Number of records. |
| records | array[file_info] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
```

```
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "accessed_time": "2019-06-12 11:00:16 -0400",
      "analytics": {
        "by_accessed_time": {
          "bytes_used": {
            "aged_data_metric": 15.23,
            "labels": [
              "2019-07",
              "2019-06",
              "2019-05",
              "2019",
              "2018",
              "--2017",
              "unknown"
            ],
            "newest_label": "2019-07",
            "oldest_label": "2019-07",
            "percentages": [
              0.1,
              11.24,
              0.18,
              15.75,
              0.75,
              83.5,
              0
            ],
            "values": [
              15925248,
              1735569408,
              27672576,
              2430595072,
              116105216,
              12889948160,
              0
            ]
          }
        },
        "by_modified_time": {
          "bytes_used": {
            "aged_data_metric": 15.23,
            "labels": [
              "2019-07",
              "2019-06",
```

```
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
          ],
          "newest_label": "2019-07",
          "oldest_label": "2019-07",
          "percentages": [
            0.1,
            11.24,
            0.18,
            15.75,
            0.75,
            83.5,
            0
          ],
          "values": [
            15925248,
            1735569408,
            27672576,
            2430595072,
            116105216,
            12889948160,
            0
          ]
        }
      },
      "bytes_used": 15436648448,
      "file_count": 21134,
      "report_time": "string",
      "subdir_count": 35
    },
    "bytes_used": 4096,
    "changed_time": "2019-06-12 11:00:16 -0400",
    "constituent": {
      "name": "fg__0001",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "creation_time": "2019-06-12 11:00:16 -0400",
    "group_id": 30,
    "hard_links_count": 1,
    "inode_generation": 214753547,
    "inode_number": 1695,
    "is_junction": "",
    "is_snapshot": "",
```

```
      "is_vm_aligned": "",
      "modified_time": "2019-06-12 11:00:16 -0400",
      "name": "string",
      "owner_id": 54738,
      "path": "string",
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string",
        "uuid": "string"
      },
      "target": "string",
      "type": "file",
      "unique_bytes": 4096,
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917752 | Volume is offline. |
| 918235 | A volume with the specified UUID was not found. |
| 6488109 | Operation not supported on FlexCache volumes. |
| 6684674 | No such file or directory. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
|---|---|---|
| labels | array[string] | Labels for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

| Name | Type | Description |
|---|---|---|
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
|---|---|---|
| labels | array[string] | Labels for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification.

| Name | Type | Description |
|---|---|---|
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

analytics

Additional file system analytics information that is invariant amongst all elements in the collection.

This property is only populated if file system analytics is enabled on the containing volume.

This analytics object captures properties that are invariant amongst all elements included in the `records` array. The invariant properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large.

| Name | Type | Description |
|---|---|---|
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |

_links

| Name | Type | Description |
|---|---|---|
| metadata | href | |
| self | href | |

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
|---|---|---|
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

analytics

Additional file system analytics information summarizing all descendants of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk |
| file_count | integer | Number of descendants |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| report_time | string | The date and time analytics information was collected. |
| subdir_count | integer | Number of sub directories |

constituent

| Name | Type | Description |
| --- | --- | --- |
| name | string | FlexGroup volume constituent name. |
| uuid | string | FlexGroup volume constituent UUID. |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

> Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request. |
| uuid | string | The unique identifier of the QoS policy. Valid in PATCH. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

file_info

Information about a single file.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| accessed_time | string | Last access time of the file in date-time format. |
| analytics | analytics | Additional file system analytics information summarizing all descendants of a directory.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection. |
| bytes_used | integer | The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range. |
| changed_time | string | Last time data or attributes changed on the file in date-time format. |
| constituent | constituent | |
| creation_time | string | Creation time of the file in date-time format. |
| fill_enabled | boolean | Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field. |
| group_id | integer | The integer ID of the group of the file owner. |
| hard_links_count | integer | The number of hard links to the file. |

| Name | Type | Description |
| --- | --- | --- |
| inode_generation | integer | Inode generation number. |
| inode_number | integer | The file inode number. |
| is_empty | boolean | Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set. |
| is_junction | boolean | Returns "true" if the directory is a junction. |
| is_snapshot | boolean | Returns "true" if the directory is a snapshot. |
| is_vm_aligned | boolean | Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume. |
| modified_time | string | Last data modification time of the file in date-time format. |
| name | string | Name of the file. |
| overwrite_enabled | boolean | Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field. |
| owner_id | integer | The integer ID of the file owner. |
| path | string | Path of the file. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".<br><br>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.<br><br>Note that a QoS policy can be set on a file, or a file's volume, but not on both. |
| size | integer | The size of the file, in bytes. |
| target | string | The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some> |
| type | string | Type of the file. |
| unique_bytes | integer | Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range. |

| Name | Type | Description |
|---|---|---|
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group. |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Write to an existing file with the supplied data**

PATCH `/storage/volumes/{volume.uuid}/files/{path}`

**Introduced In:** 9.8

Writes to an existing file with the supplied data or modifies the size, name, space reservation information, QoS policy, or hole range information of a file. Query-based PATCH operations are not supported.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| path | string | path | True | Relative path of a file in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided. |
| byte_offset | integer | query | False | Indicates the number of bytes into the file to begin writing. Use "-1" to append (default). Note that the byte-offset field is only supported for writing to a new or existing file, which requires specifying the Content-Type as 'multipart/form-data'. |
| stream_name | string | query | False | Name of stream associated with the file to write data to. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| accessed_time | string | Last access time of the file in date-time format. |

| Name | Type | Description |
| --- | --- | --- |
| analytics | analytics | Additional file system analytics information summarizing all descendants of a directory.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection. |
| bytes_used | integer | The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range. |
| changed_time | string | Last time data or attributes changed on the file in date-time format. |
| constituent | constituent | |
| creation_time | string | Creation time of the file in date-time format. |
| fill_enabled | boolean | Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field. |
| group_id | integer | The integer ID of the group of the file owner. |
| hard_links_count | integer | The number of hard links to the file. |
| inode_generation | integer | Inode generation number. |
| inode_number | integer | The file inode number. |

| Name | Type | Description |
| --- | --- | --- |
| is_empty | boolean | Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set. |
| is_junction | boolean | Returns "true" if the directory is a junction. |
| is_snapshot | boolean | Returns "true" if the directory is a snapshot. |
| is_vm_aligned | boolean | Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume. |
| modified_time | string | Last data modification time of the file in date-time format. |
| name | string | Name of the file. |
| overwrite_enabled | boolean | Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field. |
| owner_id | integer | The integer ID of the file owner. |
| path | string | Path of the file. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".<br><br>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.<br><br>Note that a QoS policy can be set on a file, or a file's volume, but not on both. |
| size | integer | The size of the file, in bytes. |
| target | string | The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some> |
| unique_bytes | integer | Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range. |

| Name | Type | Description |
| --- | --- | --- |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group. |
| volume | volume | |

**Example request**


```
{
  "accessed_time": "2019-06-12 11:00:16 -0400",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
```

```
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "bytes_used": 15436648448,
    "file_count": 21134,
    "report_time": "string",
    "subdir_count": 35
  },
  "bytes_used": 4096,
  "changed_time": "2019-06-12 11:00:16 -0400",
  "constituent": {
    "name": "fg__0001",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "creation_time": "2019-06-12 11:00:16 -0400",
  "group_id": 30,
  "hard_links_count": 1,
  "inode_generation": 214753547,
  "inode_number": 1695,
  "is_junction": "",
  "is_snapshot": "",
  "is_vm_aligned": "",
```

```
  "modified_time": "2019-06-12 11:00:16 -0400",
  "name": "string",
  "owner_id": 54738,
  "path": "string",
  "qos_policy": {
    "name": "string",
    "uuid": "string"
  },
  "target": "string",
  "unique_bytes": 4096,
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 918235 | A volume with the specified UUID was not found. |
| 6488081 | The {field} field is not supported for PATCH operations. |
| 6488082 | Failed to rename {path}. |
| 6488083 | Failed to rename {path} to {path} because a directory named {path} already exists. |
| 6488111 | The fields 'fill_enabled' and 'overwrite_enabled' are not supported in the directory modify operation. |
| 6488117 | Permission denied. |
| 6488119 | Operation not supported. |

Also see the table of common errors in the Response body overview section of this documentation.

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

analytics

Additional file system analytics information summarizing all descendants of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk |
| file_count | integer | Number of descendants |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| report_time | string | The date and time analytics information was collected. |
| subdir_count | integer | Number of sub directories |

constituent

| Name | Type | Description |
| --- | --- | --- |
| name | string | FlexGroup volume constituent name. |
| uuid | string | FlexGroup volume constituent UUID. |

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

> Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request. |
| uuid | string | The unique identifier of the QoS policy. Valid in PATCH. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

file_info

Information about a single file.

| Name | Type | Description |
| --- | --- | --- |
| accessed_time | string | Last access time of the file in date-time format. |
| analytics | analytics | Additional file system analytics information summarizing all descendants of a directory.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection. |
| bytes_used | integer | The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range. |
| changed_time | string | Last time data or attributes changed on the file in date-time format. |
| constituent | constituent | |
| creation_time | string | Creation time of the file in date-time format. |
| fill_enabled | boolean | Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field. |
| group_id | integer | The integer ID of the group of the file owner. |
| hard_links_count | integer | The number of hard links to the file. |

| Name | Type | Description |
| --- | --- | --- |
| inode_generation | integer | Inode generation number. |
| inode_number | integer | The file inode number. |
| is_empty | boolean | Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set. |
| is_junction | boolean | Returns "true" if the directory is a junction. |
| is_snapshot | boolean | Returns "true" if the directory is a snapshot. |
| is_vm_aligned | boolean | Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume. |
| modified_time | string | Last data modification time of the file in date-time format. |
| name | string | Name of the file. |
| overwrite_enabled | boolean | Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field. |
| owner_id | integer | The integer ID of the file owner. |
| path | string | Path of the file. |

| Name | Type | Description |
| --- | --- | --- |
| qos_policy | qos_policy | The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".<br><br>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.<br><br>Note that a QoS policy can be set on a file, or a file's volume, but not on both. |
| size | integer | The size of the file, in bytes. |
| target | string | The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some> |
| unique_bytes | integer | Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range. |

| Name | Type | Description |
| --- | --- | --- |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group. |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a new file with the supplied data**

POST `/storage/volumes/{volume.uuid}/files/{path}`

**Introduced In:** 9.8

Creates a new file with the supplied data, a new directory or a new symlink.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| path | string | path | True | Relative path of a new file, directory or symlink. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided. |
| byte_offset | integer | query | False | Indicates the number of bytes into the file to begin writing. Use "-1" to append (default). Note that the byte-offset field is only supported for writing to a new or existing file, which requires specifying the Content-Type as 'multipart/form-data'. |
| overwrite | boolean | query | False | If false, and the file exists, the write will fail. Default is false. |
| stream_name | string | query | False | Name of stream associated with the file to write data to. |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| accessed_time | string | Last access time of the file in date-time format. |

| Name | Type | Description |
| --- | --- | --- |
| analytics | analytics | Additional file system analytics information summarizing all descendants of a directory.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection. |
| bytes_used | integer | The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range. |
| changed_time | string | Last time data or attributes changed on the file in date-time format. |
| constituent | constituent | |
| creation_time | string | Creation time of the file in date-time format. |
| group_id | integer | The integer ID of the group of the file owner. |
| hard_links_count | integer | The number of hard links to the file. |
| inode_generation | integer | Inode generation number. |
| inode_number | integer | The file inode number. |
| is_empty | boolean | Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set. |

| Name | Type | Description |
| --- | --- | --- |
| is_junction | boolean | Returns "true" if the directory is a junction. |
| is_snapshot | boolean | Returns "true" if the directory is a snapshot. |
| is_vm_aligned | boolean | Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume. |
| modified_time | string | Last data modification time of the file in date-time format. |
| name | string | Name of the file. |
| owner_id | integer | The integer ID of the file owner. |
| path | string | Path of the file. |
| target | string | The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some> |
| type | string | Type of the file. |
| unique_bytes | integer | Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range. |

| Name | Type | Description |
| --- | --- | --- |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group. |
| volume | volume | |

**Example request**


```
{
  "accessed_time": "2019-06-12 11:00:16 -0400",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "aged_data_metric": 15.23,
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
```

```
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
          1735569408,
          27672576,
          2430595072,
          116105216,
          12889948160,
          0
        ]
      }
    },
    "bytes_used": 15436648448,
    "file_count": 21134,
    "report_time": "string",
    "subdir_count": 35
  },
  "bytes_used": 4096,
  "changed_time": "2019-06-12 11:00:16 -0400",
  "constituent": {
    "name": "fg__0001",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "creation_time": "2019-06-12 11:00:16 -0400",
  "fill_enabled": true,
  "group_id": 30,
  "hard_links_count": 1,
  "inode_generation": 214753547,
  "inode_number": 1695,
  "is_junction": "",
  "is_snapshot": "",
```

```
    "is_vm_aligned": "",
    "modified_time": "2019-06-12 11:00:16 -0400",
    "name": "string",
    "overwrite_enabled": true,
    "owner_id": 54738,
    "path": "string",
    "target": "string",
    "type": "file",
    "unique_bytes": 4096,
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
```


#### Response

```
Status: 201, Created
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 917505 | The SVM does not exist. |
| 917525 | The volume in the symlink path does not exist in the SVM. |
| 917698 | Volume (vol name) in SVM (vserver) is not mounted in the namespace. |
| 917698 | The volume in the symlink path is not mounted in the namespace. |
| 918235 | A volume with the specified UUID was not found. |
| 6488064 | This command is not supported. |

| Error Code | Description |
| --- | --- |
| 6488065 | The volume in the symlink path is invalid. |
| 6488066 | Mounting the unjunctioned volume in the symlink path failed. |
| 6488069 | Internal file error. |
| 6488084 | Failed to create {path} because the "unix_permissions" field was not specified. |
| 6488085 | Failed to create {path} because the "type" field was not specified. |
| 6488118 | File exists. |
| 8257536 | This operation is not supported for the system volume specified in the symlink path. |
| 8257541 | Failed to compute the SVM identification from this content. |
| 8257542 | This operation is not supported for the administrative SVM. |
| 9437549 | This operation is not allowed on SVMs with Infinite Volume. |
| 13172837 | This operation is not permitted because the SVM is locked for a migrate operation. |
| 111411203 | Failed to get the volume file system analytics report on the volume. |
| 111411204 | Internal error. Failed to retrieve the volume file system analytics report on the volume. |
| 111411207 | Volume file system analytics is not supported on volumes that contain LUNs. |
| 111411209 | Volume file system analytics is not supported on FlexCache volumes. |
| 111411210 | Volume file system analytics is not supported on audit staging volumes. |
| 111411211 | Volume file system analytics is not supported on object store server volumes. |
| 111411212 | Volume file system analytics is not supported on SnapMirror destination volumes. |
| 111411215 | Internal error. Volume file system analytics report timed out for volume volume.name in SVM svm.name. |
| 111411216 | Enabling or disabling volume file system analytics is not supported on individual FlexGroup constituents. |
| 111411217 | Volume file system analytics is not supported on SnapLock volumes. |

| Error Code | Description |
| --- | --- |
| 111411230 | Volume file system analytics is not supported on volumes that contain NVMe namespaces. |
| 111412203 | Volume file system analytics is not enabled on the volume. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

bytes_used

Number of bytes used on-disk, broken down by date of last access.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_accessed_time

File system analytics information, broken down by date of last access.

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last access. |

bytes_used

Number of bytes used on-disk, broken down by date of last modification.

| Name | Type | Description |
| --- | --- | --- |
| aged_data_metric | number | A score summarizing how old the data is. A higher value means the data is older. |
| labels | array[string] | Labels for this histogram |

| Name | Type | Description |
| --- | --- | --- |
| newest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| oldest_label | string | Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul><br>a partial date in an extended ISO8601 representation<br>an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval<br>the string literal "unknown"<br><br></ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.<br><br>The following extensions to ISO8601 are used:<ul><br>Quarters may be specified. The form *yyyy*-Q*q* is used to represent the *q*th quarter of the year *yyyy*. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.<br>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval |

| Name | Type | Description |
| --- | --- | --- |
| percentages | array[number] | Percentages for this histogram |
| values | array[integer] | Values for this histogram |

by_modified_time

File system analytics information, broken down by date of last modification with any other time period. This usually occurs when the data was at some point associated with a time in the future.

| Name | Type | Description |
| --- | --- | --- |
| bytes_used | bytes_used | Number of bytes used on-disk, broken down by date of last modification. |

analytics

Additional file system analytics information summarizing all descendants of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.

| Name | Type | Description |
| --- | --- | --- |
| by_accessed_time | by_accessed_time | File system analytics information, broken down by date of last access. |
| by_modified_time | by_modified_time | File system analytics information, broken down by date of last modification. |
| bytes_used | integer | Number of bytes used on-disk |
| file_count | integer | Number of descendants |
| incomplete_data | boolean | Returns true if data collection is incomplete for this directory tree. |
| report_time | string | The date and time analytics information was collected. |
| subdir_count | integer | Number of sub directories |

constituent

| Name | Type | Description |
| --- | --- | --- |
| name | string | FlexGroup volume constituent name. |
| uuid | string | FlexGroup volume constituent UUID. |

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

> Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request. |
| uuid | string | The unique identifier of the QoS policy. Valid in PATCH. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

file_info

Information about a single file.

| Name | Type | Description |
|---|---|---|
| accessed_time | string | Last access time of the file in date-time format. |
| analytics | analytics | Additional file system analytics information summarizing all descendants of a directory.<br><br>This property is only populated if file system analytics is enabled on the containing volume.<br><br>In the context of the `records` property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection. |
| bytes_used | integer | The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range. |
| changed_time | string | Last time data or attributes changed on the file in date-time format. |
| constituent | constituent | |
| creation_time | string | Creation time of the file in date-time format. |
| group_id | integer | The integer ID of the group of the file owner. |
| hard_links_count | integer | The number of hard links to the file. |
| inode_generation | integer | Inode generation number. |
| inode_number | integer | The file inode number. |

| Name | Type | Description |
| --- | --- | --- |
| is_empty | boolean | Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set. |
| is_junction | boolean | Returns "true" if the directory is a junction. |
| is_snapshot | boolean | Returns "true" if the directory is a snapshot. |
| is_vm_aligned | boolean | Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume. |
| modified_time | string | Last data modification time of the file in date-time format. |
| name | string | Name of the file. |
| owner_id | integer | The integer ID of the file owner. |
| path | string | Path of the file. |
| target | string | The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some> |
| type | string | Type of the file. |

| Name | Type | Description |
| --- | --- | --- |
| unique_bytes | integer | Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range. |
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group. |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve historical performance metrics for a volume**

GET `/storage/volumes/{volume.uuid}/metrics`

**Introduced In:** 9.7

Retrieves historical performance metrics for a volume.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| throughput.other | integer | query | False | Filter by throughput.other |
| throughput.total | integer | query | False | Filter by throughput.total |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| timestamp | string | query | False | Filter by timestamp |
| iops.other | integer | query | False | Filter by iops.other |
| iops.total | integer | query | False | Filter by iops.total |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| latency.other | integer | query | False | Filter by latency.other |
| latency.total | integer | query | False | Filter by latency.total |
| latency.write | integer | query | False | Filter by latency.write |
| latency.read | integer | query | False | Filter by latency.read |
| status | string | query | False | Filter by status |
| cloud.duration | string | query | False | Filter by cloud.duration |
| cloud.status | string | query | False | Filter by cloud.status |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| cloud.timestamp | string | query | False | Filter by cloud.timestamp |
| cloud.iops.other | integer | query | False | Filter by cloud.iops.other |
| cloud.iops.total | integer | query | False | Filter by cloud.iops.total |
| cloud.iops.write | integer | query | False | Filter by cloud.iops.write |
| cloud.iops.read | integer | query | False | Filter by cloud.iops.read |
| cloud.latency.other | integer | query | False | Filter by cloud.latency.other |
| cloud.latency.total | integer | query | False | Filter by cloud.latency.total |
| cloud.latency.write | integer | query | False | Filter by cloud.latency.write |
| cloud.latency.read | integer | query | False | Filter by cloud.latency.read |
| flexcache.timestamp | string | query | False | Filter by flexcache.timestamp<br>• Introduced in: 9.8 |
| flexcache.cache_miss_percent | integer | query | False | Filter by flexcache.cache_miss_percent<br>• Introduced in: 9.8 |
| flexcache.bandwidth_savings | integer | query | False | Filter by flexcache.bandwidth_savings<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| flexcache.duration | string | query | False | Filter by flexcache.duration<br><br>• Introduced in: 9.8 |
| flexcache.status | string | query | False | Filter by flexcache.status<br><br>• Introduced in: 9.8 |
| duration | string | query | False | Filter by duration |
| volume.uuid | string | path | True | Unique identifier of the volume. |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| interval | string | query | False | The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.<br>The period for each time range is as follows:<br>• 1h: Metrics over the most recent hour sampled over 15 seconds.<br>• 1d: Metrics over the most recent day sampled over 5 minutes.<br>• 1w: Metrics over the most recent week sampled over 30 minutes.<br>• 1m: Metrics over the most recent month sampled over 2 hours.<br>• 1y: Metrics over the most recent year sampled over a day.<br>• Default value: 1<br>• enum: ["1h", "1d", "1w", "1m", "1y"] |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | return_records | boolean | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 +0000"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

iops

The rate of I/O operations observed at the storage object.

| Name | Type | Description |
|---|---|---|
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

throughput

The rate of throughput bytes per second observed at the storage object.

| Name | Type | Description |
| --- | --- | --- |
| other | integer | Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on. |
| read | integer | Performance metric for read I/O operations. |
| total | integer | Performance metric aggregated over all types of I/O operations. |
| write | integer | Performance metric for write I/O operations. |

records

Performance numbers, such as IOPS latency and throughput.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage volume snapshots

### Manage volume snapshots

#### Overview

A snapshot is the view of the filesystem as it exists at the time when the snapshot is created.

In ONTAP, different types of snapshots are supported, such as scheduled snapshots, user requested snapshots, SnapMirror snapshots, and so on.

ONTAP snapshot APIs allow you to create, modify, delete and retrieve snapshots.

ONTAP Bulk snapshot APIs allow you to create, modify, delete and retrieve snapshots on multiple volumes in one request.

#### Snapshot APIs

The following APIs are used to perform operations related to snapshots.

– POST /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– PATCH /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– DELETE /api/storage/volumes/{volume.uuid}/snapshots/{uuid}<br> The following APIs are used to perform bulk operations related to snapshots.

– POST /api/storage/volumes/*/snapshots

– GET /api/storage/volumes/*/snapshots

– PATCH /api/storage/volumes/*/snapshots/{uuid}

– DELETE /api/storage/volumes/*/snapshots/{uuid}

#### Examples

### Creating a snapshot

The POST operation is used to create a snapshot with the specified attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X POST "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json" -d '{"name": "snapshot_copy", "comment": "Store this
copy." }'

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:43:34 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy
Content-Length: 189
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "volume": {
      "name": "v2"
    },
    "svm": {
      "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
      "name": "vs0"
    },
    "name": "snapshot_copy",
    "comment": "Store this copy."
  }
],
```

```
"job": {
  "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6f68c85b-45e1-11e9-8fc7-005056bbc848"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:43:57 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 224
Content-Type: application/json
{
"uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
"description": "POST /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy",
"state": "success",
"message": "success",
"code": 0
}
```

**Retrieving snapshot attributes**

The GET operation is used to retrieve snapshot attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 21:14:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/json
```

```
Transfer-Encoding: chunked
{
"records": [
  {
    "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
    "name": "hourly.2019-03-13_1305",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
      }
    }
  },
  {
    "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
    "name": "hourly.2019-03-13_1405",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
      }
    }
  },
  {
    "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
    "name": "hourly.2019-03-13_1522",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots"
  }
}
}
```

**Creating bulk snapshots**

The POST operation is used to create a snapshot with the same name on multiple volumes in one request.

This operation accepts a volume UUID or volume name and SVM, and a snapshot name.
This operation only supports SnapMirror label attributes to be added to snapshots during creation.

```
# The API:
/api/storage/volumes/*/snapshots

# The call:
curl -k -u admin -X POST "https://<mgmt-
ip>/api/storage/volumes/*/snapshots" -H 'accept: application/hal+json' -d
'{"records": [{"volume.uuid":"e8815adb-5209-11ec-b4ad-005056bbc3e8",
"name":"snapshot_copy"},{ "volume.uuid":"efda9101-5209-11ec-b4ad-
005056bbc3e8", "name":"snapshot_copy"}]}'

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 14 Dec 2021 20:18:13 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/storage/volumes/*/snapshots/
Content-Length: 209
Content-Type: application/json
{
"num_records": 2,
"job": {
  "uuid": "f7130fc0-5d1a-11ec-b78c-005056bbb467",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f7130fc0-5d1a-11ec-b78c-005056bbb467"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Tue, 14 Dec 2021 20:20:54 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
```

```
Content-Length: 258
Content-Type: application/json
Vary: Accept-Encoding
{
"uuid": "f7130fc0-5d1a-11ec-b78c-005056bbb467",
"description": "POST /api/storage/volumes/*/snapshots/",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-12-14T15:18:13-05:00",
"end_time": "2021-12-14T15:18:13-05:00"
}
```

**Retrieving snapshot advanced attributes**

A collection GET request is used to calculate the amount of snapshot reclaimable space.
When the advanced privilege field 'reclaimable space' is requested, the API returns the amount of reclaimable space for the queried list of snapshots.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots?fields=reclaimable_space

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=reclaimable_space&n
ame=hourly.2019-03-13_1305&#124;hourly.2019-03-13_1405&#124;hourly.2019-
03-13_1522" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 21:14:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/json
Transfer-Encoding: chunked
{
"records": [
  {
    "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
    "name": "hourly.2019-03-13_1305",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
      }
```

```
    }
  },
  {
    "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
    "name": "hourly.2019-03-13_1405",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
      }
    }
  },
  {
    "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
    "name": "hourly.2019-03-13_1522",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
      }
    }
  }
],
"num_records": 3,
"reclaimable_space": 1567832,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots"
  }
}
}
```


**Retrieving snapshot advanced attributes**

A collection GET request is used to calculate the delta between two snapshots.
When the advanced privilege field 'delta' is requested, the API returns the delta between the queried snapshots.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots?fields=delta

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=delta&name=hourly.2
022-06-29_1105,hourly.2022-06-29_1205" -H "accept: application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
Connection: Keep-Alive
Content-Encoding: gzip
Content-Length: 378
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Type: application/hal+json
Date: Wed, 29 Jun 2022 18:37:11 GMT
Keep-Alive: timeout=5, max=99
Server: libzapid-httpd
Vary: Accept-Encoding
X-Content-Type-Options: nosniff
{
"records": [
  {
    "uuid": "52a2247a-7735-4a92-bc3c-e51df1fe502f",
    "name": "hourly.2022-06-29_1105",
    "delta": {
      "size_consumed": 675840,
      "time_elapsed": "PT3H27M45S"
    }
  },
  {
    "uuid": "b399eb34-44fe-4689-9fb5-c8f72162dd77",
    "name": "hourly.2022-06-29_1205",
    "delta": {
      "size_consumed": 507904,
      "time_elapsed": "PT2H27M45S"
    }
  }
],
"num_records": 2,
"delta": {
  "size_consumed": 167936,
  "time_elapsed": "PT1H"
}
}
```

**Retrieving the attributes of a specific snapshot**

The GET operation is used to retrieve the attributes of a specific snapshot.

```
# The API:
```

```
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:39:26 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 308
Content-Type: application/json
{
"volume": {
  "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",
  "name": "v2",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"
    }
  }
},
"uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
"svm": {
  "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"
    }
  }
},
"name": "hourly.2019-03-13_1305",
"create_time": "2019-03-13T13:05:00-04:00",
"size": 122880,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
  }
}
}
```

#### Retrieving the advanced attributes of a specific snapshot

The GET operation is used to retrieve the attributes of a specific snapshot. Snapshot reclaimable space can be requested during a GET request.
When the advanced privilege field reclaimable space is requested, the API returns the amount of reclaimable space for the snapshot.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}?fields=reclaimable_spa
ce

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-
75ee0ab3e8c0?fields=**" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:39:26 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 308
Content-Type: application/json
{
"volume": {
  "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",
  "name": "v2",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"
    }
  }
},
"uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
"svm": {
  "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"
    }
  }
},
"name": "hourly.2019-03-13_1305",
"reclaimable_space": 167832,
"_links": {
```

```
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
  }
}
}
```

**Retrieving snapshot advanced attributes**

A collection GET request is used to calculate the delta between two snapshots.
When the advanced privilege field 'delta' is requested, the API returns the delta between the queried snapshots.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots?fields=delta

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=delta&name=hourly.2
022-06-29_1105,hourly.2022-06-29_1205" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Connection: Keep-Alive
Content-Encoding: gzip
Content-Length: 378
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Type: application/hal+json
Date: Wed, 29 Jun 2022 18:37:11 GMT
Keep-Alive: timeout=5, max=99
Server: libzapid-httpd
Vary: Accept-Encoding
X-Content-Type-Options: nosniff
{
"records": [
  {
    "uuid": "52a2247a-7735-4a92-bc3c-e51df1fe502f",
    "name": "hourly.2022-06-29_1105",
    "delta": {
      "size_consumed": 675840,
      "time_elapsed": "PT3H27M45S"
    }
  },
  {
```

```
    "uuid": "b399eb34-44fe-4689-9fb5-c8f72162dd77",
    "name": "hourly.2022-06-29_1205",
    "delta": {
      "size_consumed": 507904,
      "time_elapsed": "PT2H27M45S"
    }
  }
],
"num_records": 2,
"delta": {
  "size_consumed": 167936,
  "time_elapsed": "PT1H"
}
}
```

**Retrieving bulk snapshots**

The bulk GET operation is used to retrieve snapshot attributes across all volumes.

```
# The API:
/api/storage/volumes/*/snapshots

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
{
"records": [
  {
    "volume": {
      "uuid": "966c285f-47f7-11ec-8407-005056bbc08f",
      "name": "v1"
    },
    "uuid": "3edba912-5507-4535-adce-e12fe5c0e31c",
    "name": "daily.2021-11-18_0010"
  },
  {
    "volume": {
      "uuid": "966c285f-47f7-11ec-8407-005056bbc08f",
      "name": "v1"
    },
    "uuid": "3ad61153-d5ef-495d-8e0e-5c3b8bbaf5e6",
    "name": "hourly.2021-11-18_0705"
  },
```

```
    {
      "volume": {
        "uuid": "99c974e3-47f7-11ec-8407-005056bbc08f",
        "name": "v2"
      },
      "uuid": "3dd0fa97-65d9-41ea-a99d-5ceb9d2f55c5",
      "name": "daily.2021-11-18_0010"
    },
    {
     "volume": {
       "uuid": "99c974e3-47f7-11ec-8407-005056bbc08f",
       "name": "v2"
     },
     "uuid": "6ca20a52-c342-4753-8865-3693fa9b7e23",
     "name": "hourly.2021-11-18_0705"
    },
  ],
  "num_records": 4
  }
```


**Updating a snapshot**

The PATCH operation is used to update the specific attributes of a snapshot.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -d
'{"name": "snapshot_copy_new" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:50:44 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6f7c3a82-45e2-11e9-8fc7-005056bbc848"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:54:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 242
Content-Type: application/json
{
"uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
"description": "PATCH /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
"state": "success",
"message": "success",
"code": 0
}
```

### Updating bulk snapshots

The bulk PATCH operation is used to update the specific attributes of snapshots across volumes in a single request.

```
# The API:
/api/storage/volumes/*/snapshots

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -d
'{"records": [{"volume.uuid":"e8815adb-5209-11ec-b4ad-005056bbc3e8",
"svm.uuid":"d0e6def5-5209-11ec-b4ad-005056bbc3e8", "uuid":"f9b7714d-1166-
410a-b143-874f27969db6", "comment":"yay"},{"volume.uuid":"efda9101-5209-
11ec-b4ad-005056bbc3e8", "svm.uuid":"d0e6def5-5209-11ec-b4ad-
005056bbc3e8", "uuid":"514c82a7-bff7-48e2-a13c-5337b09ed41e",
"comment":"yay"}]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
{
"job": {
  "uuid": "1e9a561f-520f-11ec-b4ad-005056bbc3e8",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/1e9a561f-520f-11ec-b4ad-005056bbc3e8"
    },
    "results": {
      "href": "/api/storage/volumes/*/snapshots?job_results_uuid=1e9a561f-
520f-11ec-b4ad-005056bbc3e8"
    }
  }
}
}

# The Job:
curl -u admin:<password> -k -X GET --header 'Content-Type:
application/json' --header 'Accept: application/json' 'https://<mgmt-
ip>/api/storage/volumes/*/snapshots?job_results_uuid=1e9a561f-520f-11ec-
b4ad-005056bbc3e8'
HTTP/1.1 200 OK
{
"records": [
  {
    "volume": {
      "uuid": "e8815adb-5209-11ec-b4ad-005056bbc3e8",
      "name": "v1"
    },
```

```
    "uuid": "f9b7714d-1166-410a-b143-874f27969db6",
    "svm": {
      "uuid": "d0e6def5-5209-11ec-b4ad-005056bbc3e8"
    },
    "name": "s1",
    "comment": "yay"
  },
  {
    "volume": {
      "uuid": "efda9101-5209-11ec-b4ad-005056bbc3e8",
      "name": "v2"
    },
    "uuid": "514c82a7-bff7-48e2-a13c-5337b09ed41e",
    "svm": {
      "uuid": "d0e6def5-5209-11ec-b4ad-005056bbc3e8"
    },
    "name": "s1",
    "comment": "yay"
  }
],
"num_records": 2
}
```


**Deleting a snapshot**

The DELETE operation is used to delete a snapshot.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:57:51 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6da1dfdd-45e3-11e9-8fc7-005056bbc848"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 23:02:46 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 243
Content-Type: application/json
{
"uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
"description": "DELETE /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
"state": "success",
"message": "success",
"code": 0
}
```

#### Deleting bulk snapshots

The bulk DELETE operation is used to delete a snapshots across volumes in a single request.

```
# The API:
/api/storage/volumes/*/snapshots

# The call:
curl -X DELETE  "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -d
'{"records": [{"volume.uuid":"e8815adb-5209-11ec-b4ad-005056bbc3e8",
"uuid":"f9b7714d-1166-410a-b143-874f27969db6"}, {"volume.uuid":"efda9101-
5209-11ec-b4ad-005056bbc3e8", "uuid":"1d55c97a-25f3-4366-bfa8-
9ea75c255469"}]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
{
"job": {
  "uuid": "fe114ed7-520f-11ec-b4ad-005056bbc3e8",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/fe114ed7-520f-11ec-b4ad-005056bbc3e8"
    },
    "results": {
      "href": "/api/storage/volumes/*/snapshots?job_results_uuid=fe114ed7-
520f-11ec-b4ad-005056bbc3e8"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
curl -u admin:<password> -k -X GET --header 'Content-Type:
application/json' --header 'Accept: application/json' 'https://<mgmt-
ip>/api/storage/volumes/*/snapshots?job_results_uuid=fe114ed7-520f-11ec-
b4ad-005056bbc3e8'
{
"records": [
],
"num_records": 0
}
```

#### Retrieve volume snapshots

GET `/storage/volumes/{volume.uuid}/snapshots`

**Introduced In:** 9.6

Retrieves a collection of volume snapshots.

#### Expensive properties

There is an added computational cost to retrieving the amount of reclaimable space for snapshots, as the calculation is done on demand based on the list of snapshots provided.

- `reclaimable_space`
- `delta`

#### Related ONTAP commands

- `snapshot show`
- `snapshot compute-reclaimable`
- `snapshot show-delta`

#### Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume |
| version_uuid | string | query | False | Filter by version_uuid<br>• Introduced in: 9.11 |
| comment | string | query | False | Filter by comment |
| owners | string | query | False | Filter by owners<br>• Introduced in: 9.7 |
| volume.name | string | query | False | Filter by volume.name |
| uuid | string | query | False | Filter by uuid |
| size | integer | query | False | Filter by size<br>• Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |
| state | string | query | False | Filter by state |
| dedup_savings | integer | query | False | Filter by dedup_savings<br>• Introduced in: 9.17 |
| snapmirror_label | string | query | False | Filter by snapmirror_label<br>• Introduced in: 9.8 |
| delta.size_consumed | integer | query | False | Filter by delta.size_consumed<br>• Introduced in: 9.12 |
| delta.time_elapsed | string | query | False | Filter by delta.time_elapsed<br>• Introduced in: 9.12 |
| reclaimable_space | integer | query | False | Filter by reclaimable_space<br>• Introduced in: 9.10 |
| create_time | string | query | False | Filter by create_time |
| expiry_time | string | query | False | Filter by expiry_time |
| compress_savings | integer | query | False | Filter by compress_savings<br>• Introduced in: 9.17 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| vbn0_savings | integer | query | False | Filter by vbn0_savings<br>• Introduced in: 9.17 |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| provenance_volume.uuid | string | query | False | Filter by provenance_volume.uuid<br>• Introduced in: 9.11 |
| logical_size | integer | query | False | Filter by logical_size<br>• Introduced in: 9.12 |
| snaplock_expiry_time | string | query | False | Filter by snaplock_expiry_time |
| snaplock.expiry_time | string | query | False | Filter by snaplock.expiry_time<br>• Introduced in: 9.15 |
| snaplock.time_until_expiry | string | query | False | Filter by snaplock.time_until_expiry<br>• Introduced in: 9.15 |
| snaplock.expired | boolean | query | False | Filter by snaplock.expired<br>• Introduced in: 9.15 |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |
| num_records | integer | Number of records |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| records | array[snapshot] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "delta": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "compress_savings": 1131223,
      "create_time": "2019-02-04 19:00:00 +0000",
      "dedup_savings": 1131223,
      "delta": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "size_consumed": 0,
        "time_elapsed": "string"
      },
      "expiry_time": "2019-02-04 19:00:00 +0000",
      "logical_size": 1228800,
      "name": "this_snapshot",
      "owners": [
        "string"
```

```
      ],
      "provenance_volume": {
        "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
      },
      "size": "122880",
      "snaplock": {
        "expired": 1,
        "expiry_time": "2019-02-04 19:00:00 +0000",
        "time_until_expiry": "PT3H27M45S"
      },
      "snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
      "snapmirror_label": "string",
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "vbn0_savings": 1131223,
      "version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
| --- | --- |
| 918235 | The specified volume is invalid. |
| 9437613 | The operation is not supported on FlexCache Volumes. |

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

snapshot_delta

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| size_consumed | integer | Indicates the space that has changed between two specified WAFL file systems, in bytes. |
| time_elapsed | string | Time elapsed between two specified WAFL file systems. |

provenance_volume

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| expired | boolean | Indicates whether a SnapLock snapshot has expired. |
| expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. It can also be used to extend the expiry time of a locked snapshot on a SnapLock for SnapVault destination consistency-group. |
| time_until_expiry | string | Indicates the remaining SnapLock expiry time of a locked snapshot, in seconds. This field is set only when the remaining time interval is less than 136 years. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

snapshot

The snapshot object represents a point in time snapshot of a volume.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |

| Name | Type | Description |
| --- | --- | --- |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |

| Name | Type | Description |
| --- | --- | --- |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |
| snapmirror_label | string | Label for SnapMirror operations |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |

| Name | Type | Description |
|---|---|---|
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create a volume snapshot

POST `/storage/volumes/{volume.uuid}/snapshots`

**Introduced In:** 9.6

Creates a volume snapshot.

#### Required properties

- `name` - Name of the snapshot to be created.

#### Recommended optional properties

- `comment` - Comment associated with the snapshot.
- `expiry_time` - snapshots with an expiry time set are not allowed to be deleted until the retention time is reached.
- `snapmirror_label` - Label for SnapMirror operations.
- `snaplock_expiry_time` - Expiry time for snapshot locking enabled volumes.

#### Related ONTAP commands

- `snapshot create`

#### Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |
| snapmirror_label | string | Label for SnapMirror operations |

| Name | Type | Description |
| --- | --- | --- |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |
| volume | volume | |

**Example request**


```
{
  "comment": "string",
  "compress_savings": 1131223,
  "create_time": "2019-02-04 19:00:00 +0000",
  "dedup_savings": 1131223,
  "delta": {
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "expiry_time": "2019-02-04 19:00:00 +0000",
  "logical_size": 1228800,
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "provenance_volume": {
    "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
  },
  "size": "122880",
  "snaplock": {
    "expired": 1,
    "expiry_time": "2019-02-04 19:00:00 +0000",
    "time_until_expiry": "PT3H27M45S"
  },
  "snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
  "snapmirror_label": "string",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vbn0_savings": 1131223,
  "version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
|---|---|
| 524479 | The specified volume is not online or does not have enough space to create a snapshot. |
| 1520351 | Snapshot operation is not allowed. |
| 525062 | Cannot exceed maximum number of snapshots. |
| 2621462 | The specified SVM name does not exist. |
| 525059 | A snapshot with the specified name already exists. |

| Error Code | Description |
| --- | --- |
| 1638461 | Snapshots can only be created on read/write (RW) volumes. |
| 1638477 | User-created snapshot names cannot begin with the specified prefix. |
| 1638518 | The specified snapshot name is invalid. |
| 1638532 | Failed to create the snapshot on the specified volume because a revert operation is in progress. |
| 1638537 | Cannot determine the status of the snapshot create operation for the specified volume. |
| 1638574 | The "expiry-time" parameter is not supported for FlexGroup volumes. |
| 1638616 | Bulk snapshot create is not supported with multiple snapshot names. |
| 1638617 | Bulk snapshot create is not supported with volume names in a mixed-version cluster. |
| 1638618 | The property cannot be specified for snapshot create. |
| 1638630 | Snapshot create operation is not supported on destination of SnapMirror active sync relationship with policy-type "automated-failover-duplex". |
| 1638644 | POST, DELETE, and PATCH requests on the snapshot session endpoint are not supported on this platform. |
| 92405873 | Unable to create the snapshot for object store volumes. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

snapshot_delta

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned.

| Name | Type | Description |
|---|---|---|
| size_consumed | integer | Indicates the space that has changed between two specified WAFL file systems, in bytes. |
| time_elapsed | string | Time elapsed between two specified WAFL file systems. |

provenance_volume

| Name | Type | Description |
|---|---|---|
| uuid | string | UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only. |

snaplock

| Name | Type | Description |
|---|---|---|
| expired | boolean | Indicates whether a SnapLock snapshot has expired. |

| Name | Type | Description |
| --- | --- | --- |
| expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. It can also be used to extend the expiry time of a locked snapshot on a SnapLock for SnapVault destination consistency-group. |
| time_until_expiry | string | Indicates the remaining SnapLock expiry time of a locked snapshot, in seconds. This field is set only when the remaining time interval is less than 136 years. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

snapshot

The snapshot object represents a point in time snapshot of a volume.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |

| Name | Type | Description |
| --- | --- | --- |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |

| Name | Type | Description |
| --- | --- | --- |
| snapmirror_label | string | Label for SnapMirror operations |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |
| volume | volume | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a volume snapshot

DELETE `/storage/volumes/{volume.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Deletes a Volume snapshot.

#### Related ONTAP commands

- `snapshot delete`

#### Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| uuid | string | path | True | Snapshot UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
|---|---|
| 2 | An invalid value was entered for one of the fields. |
| 524481 | Snapshot was not deleted because the associated volume is locked or fenced. |
| 1638521 | Snapshots can only be deleted on read/write (RW) volumes. |
| 1638538 | Cannot determine the status of the snapshot delete operation for the specified volume. |
| 1638543 | Failed to delete snapshot because it has an owner. |
| 1638555 | The specified snapshot has not expired or is locked. |
| 1638600 | The snapshot does not exist. |
| 1638644 | POST, DELETE, and PATCH requests on the snapshot session endpoint are not supported on this platform. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

job_link

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve volume snapshot details**

GET `/storage/volumes/{volume.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Retrieves details of a specific volume snapshot.

#### Related ONTAP commands

- `snapshot show`

#### Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| uuid | string | path | True | Snapshot UUID |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |

| Name | Type | Description |
| --- | --- | --- |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |
| snapmirror_label | string | Label for SnapMirror operations |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |

| Name | Type | Description |
| --- | --- | --- |
| volume | volume | |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "compress_savings": 1131223,
  "create_time": "2019-02-04 19:00:00 +0000",
  "dedup_savings": 1131223,
  "delta": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "expiry_time": "2019-02-04 19:00:00 +0000",
  "logical_size": 1228800,
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "provenance_volume": {
    "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
  },
  "size": "122880",
  "snaplock": {
    "expired": 1,
    "expiry_time": "2019-02-04 19:00:00 +0000",
    "time_until_expiry": "PT3H27M45S"
  },
  "snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
  "snapmirror_label": "string",
  "state": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vbn0_savings": 1131223,
  "version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
|---|---|
| 2 | An invalid value was entered for one of the fields. |
| 262197 | An invalid field was specified in the request. |
| 1638473 | Snapshot tag not found. |
| 1638503 | The snapshot does not exist on the specified volume. |
| 1638600 | The snapshot does not exist. |
| 1638615 | Bulk operations for snapshots are not supported on multiple SVMs. |

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

snapshot_delta

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| size_consumed | integer | Indicates the space that has changed between two specified WAFL file systems, in bytes. |
| time_elapsed | string | Time elapsed between two specified WAFL file systems. |

provenance_volume

| Name | Type | Description |
|---|---|---|
| uuid | string | UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only. |

snaplock

| Name | Type | Description |
| --- | --- | --- |
| expired | boolean | Indicates whether a SnapLock snapshot has expired. |
| expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. It can also be used to extend the expiry time of a locked snapshot on a SnapLock for SnapVault destination consistency-group. |
| time_until_expiry | string | Indicates the remaining SnapLock expiry time of a locked snapshot, in seconds. This field is set only when the remaining time interval is less than 136 years. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update a volume snapshot**

PATCH `/storage/volumes/{volume.uuid}/snapshots/{uuid}`

**Introduced In:** 9.6

Updates a Volume snapshot.

#### Related ONTAP commands

- `snapshot modify`
- `snapshot rename`

#### Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| uuid | string | path | True | Snapshot UUID |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |

| Name | Type | Description |
| --- | --- | --- |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |
| snapmirror_label | string | Label for SnapMirror operations |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |

| Name | Type | Description |
| --- | --- | --- |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |
| volume | volume | |

**Example request**


```
{
  "comment": "string",
  "compress_savings": 1131223,
  "create_time": "2019-02-04 19:00:00 +0000",
  "dedup_savings": 1131223,
  "delta": {
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "expiry_time": "2019-02-04 19:00:00 +0000",
  "logical_size": 1228800,
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "provenance_volume": {
    "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
  },
  "size": "122880",
  "snaplock": {
    "expired": 1,
    "expiry_time": "2019-02-04 19:00:00 +0000",
    "time_until_expiry": "PT3H27M45S"
  },
  "snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
  "snapmirror_label": "string",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vbn0_savings": 1131223,
  "version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

#### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

#### Response

```
Status: 202, Accepted
```

#### Error

```
Status: Default
```

ONTAP Error Response Code

| Error Code | Description |
|---|---|
| 524508 | The snapshot was not renamed because the name entered is not valid. |
| 542797 | The specified file or snapshot does not exist. |
| 1638455 | Failed to set comment for snapshot. |
| 1638476 | You cannot rename a snapshot created for use as a reference snapshot by other jobs. |
| 1638477 | User-created snapshot names cannot begin with the specified prefix. |
| 1638518 | The specified snapshot name is invalid. |
| 1638522 | Snapshots can only be renamed on read/write (RW) volumes. |
| 1638523 | Failed to set the specified SnapMirror label for the snapshot. |

| Error Code | Description |
| --- | --- |
| 1638524 | Adding SnapMirror labels is not allowed in a mixed version cluster. |
| 1638539 | Cannot determine the status of the snapshot rename operation for the specified volume. |
| 1638554 | Failed to set expiry time for the snapshot. |
| 1638600 | The snapshot does not exist. |
| 1638644 | POST, DELETE, and PATCH requests on the snapshot session endpoint are not supported on this platform. |

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

snapshot_delta

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned.

| Name | Type | Description |
|---|---|---|
| size_consumed | integer | Indicates the space that has changed between two specified WAFL file systems, in bytes. |
| time_elapsed | string | Time elapsed between two specified WAFL file systems. |

provenance_volume

| Name | Type | Description |
|---|---|---|
| uuid | string | UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only. |

snaplock

| Name | Type | Description |
|---|---|---|
| expired | boolean | Indicates whether a SnapLock snapshot has expired. |

| Name | Type | Description |
| --- | --- | --- |
| expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. It can also be used to extend the expiry time of a locked snapshot on a SnapLock for SnapVault destination consistency-group. |
| time_until_expiry | string | Indicates the remaining SnapLock expiry time of a locked snapshot, in seconds. This field is set only when the remaining time interval is less than 136 years. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

volume

| Name | Type | Description |
| --- | --- | --- |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

snapshot

The snapshot object represents a point in time snapshot of a volume.

| Name | Type | Description |
| --- | --- | --- |
| comment | string | A comment associated with the snapshot. This is an optional attribute for POST or PATCH. |
| compress_savings | integer | Savings due to compression at the time the snapshot was taken in bytes. |
| create_time | string | Creation time of the snapshot. It is the volume access time when the snapshot was created. |
| dedup_savings | integer | Savings due to dedup at the time the snapshot was taken in bytes. |
| delta | snapshot_delta | Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a snapshot and the Active File System, only the snapshot name needs to be mentioned. |

| Name | Type | Description |
| --- | --- | --- |
| expiry_time | string | The expiry time for the snapshot. This is an optional attribute for POST or PATCH. Snapshots with an expiry time set are not allowed to be deleted until the retention time is reached. |
| logical_size | integer | Size of the logical used file system at the time the snapshot is captured. |
| name | string | Snapshot. Valid in POST or PATCH. |
| owners | array[string] | |
| provenance_volume | provenance_volume | |
| reclaimable_space | integer | Space reclaimed when the snapshot is deleted, in bytes. |
| size | integer | Size of the active file system at the time the snapshot is captured. The actual size of the snapshot also includes those blocks trapped by other snapshots. On a snapshot deletion, the "size" amount of blocks is the maximum number of blocks available. On a snapshot restore, the "afs-used size" value will match the snapshot "size" value. |
| snaplock | snaplock | |
| snaplock_expiry_time | string | SnapLock expiry time for the snapshot, if the snapshot is taken on a SnapLock volume. A snapshot is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during snapshot POST and snapshot PATCH on snapshot locking enabled volumes. This field will no longer be supported in a future release. Use snaplock.expiry_time instead. |

| Name | Type | Description |
| --- | --- | --- |
| snapmirror_label | string | Label for SnapMirror operations |
| state | string | State of the FlexGroup volume snapshot. In the "pre_conversion" state, the snapshot was created before converting the FlexVol to a FlexGroup volume. A recently created snapshot can be in the "unknown" state while the system is calculating the state. In the "partial" state, the snapshot is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial snapshots cannot be used for a snapshot restore operation. A snapshot is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a snapshot is valid. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | The UUID of the snapshot in the volume that uniquely identifies the snapshot in that volume. |
| vbn0_savings | integer | Savings due vbn0 at the time the snapshot was taken in bytes. |
| version_uuid | string | The 128 bit identifier that uniquely identifies a snapshot and its logical data layout. |
| volume | volume | |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View I/O activity for clients

### View I/O activity for clients

#### Overview

You can use this API to retrieve a list of clients with the most I/O activity for a specified volume, within the past several seconds. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` and `throughput` metrics. This API supports only returning one I/O activity type per request.

#### Approximate accounting and error bars

When too many clients have recent activity, some clients might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the clients that would have otherwise been included might not be present in the final list returned by the API.

#### Failure to return list of clients with most I/O activity

The API can sometimes fail to return the list of clients with the most I/O activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume has not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of clients with the most

I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Retrieve a list of the clients with the most I/O activity

For a report on the clients with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` I/O activity type into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of clients with the greatest number of average read operations per second. The current maximum number of clients returned by the API for an I/O activity type is 25.

– GET /api/storage/volumes/{volume.uuid}/top-metrics/clients

#### Examples

**Retrieving a list of the clients with the greatest average number of write operations per second:**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/clients

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/clients?top_metric=iops.write"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "iops": {
      "write": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "client_ip": "172.28.71.128",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
```

```
    },
    {
      "volume": {
        "name": "vol1"
      },
      "iops": {
        "write": 1022,
        "error": {
          "lower_bound": 1022,
          "upper_bound": 1032
        }
      },
      "client_ip": "172.28.71.179",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
          }
        }
      }
    },
    {
      "volume": {
        "name": "vol1"
      },
      "iops": {
        "write": 345,
        "error": {
          "lower_bound": 345,
          "upper_bound": 355
        }
      },
      "client_ip": "172.28.51.62",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
          }
        }
      }
    }
  ],
```

```
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/clients?top_metric=iops.write"
  }
}
}
```

**Example showing the behavior of the API when there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/clients?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/clients?top_metric=throughput.write"
  }
}
}
```

### Retrieve clients with the most I/O activity

GET `/storage/volumes/{volume.uuid}/top-metrics/clients`

**Introduced In:** 9.10

Retrieves a list of clients with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| volume.uuid | string | path | True | Volume UUID |
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| client_ip | string | query | False | Filter by client_ip |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| iops.write | integer | query | False | Filter by iops.write |
| iops.read | integer | query | False | Filter by iops.read |
| volume.name | string | query | False | Filter by volume.name |
| svm.name | string | query | False | Filter by svm.name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| svm.uuid | string | query | False | Filter by svm.uuid |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| incomplete_response_reason | incomplete_response_reason | Indicates that the metric report provides incomplete data. |
| notice | notice | Optional field that indicates why no records are returned by the volume activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_client] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: The activity tracking report for this volume is not available
because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume"
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "client_ip": "192.168.185.170",
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 5,
        "write": 10
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
```

```
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "throughput": {
          "error": {
            "lower_bound": 34,
            "upper_bound": 54
          },
          "read": 12,
          "write": 2
        },
        "volume": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |

| Error Code | Description |
|---|---|
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518415 | Failed to get the volume activity tracking report on volume volume.name in SVM svm.name. |
| 124518416 | Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name. |
| 124518417 | Volume wildcard queries are not supported for activity tracking reports. |
| 124518418 | The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic. Refer to the REST API documentation for more information. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 124518423 | Volume activity tracking is not supported on the volume. |
| 124519411 | Volume activity tracking is not enabled on the volume. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

incomplete_response_reason

Indicates that the metric report provides incomplete data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

notice

Optional field that indicates why no records are returned by the volume activity tracking REST API.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

top_metric_value_error_bounds

| Name | Type | Description |
|---|---|---|
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_client

Information about a client's IO activity.

| Name | Type | Description |
| --- | --- | --- |
| client_ip | string | IP address of the client. Both IPv4 and IPv6 IP addresses are supported. |
| iops | iops | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage directories

### Manage directories

#### Overview

You can use this API to retrieve a list of directories with the greatest value performance metric or capacity metric for a specified volume. Use the `top_metric` parameter to specify which type of metric to filter for. This API is used to provide insight into performance metrics, namely `iops` and `throughput`, or capacity metrics such as `non_recursive_bytes_used`. This API only supports returning one metric per request.

#### Retrieve a list of the directories with the most I/O activity

For a report on the directories with the most I/O activity in the past several seconds returned in descending order, specify the performance metric type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the metric type is not specified, by default the API returns a list of the directories with the greatest number of the average read operations per second. The current maximum number of directories returned by the API for a metric type is 25.

#### Approximate accounting and error bars for the list of directories with the most I/O activity

When too many directories have recent activity, some directories might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the directories that would have otherwise been included might not be present in the final list returned by the API.

#### Retrieve a list of the largest directories

For a report on the largest directories returned in descending order, specify the capacity metric by passing the `non_recursive_bytes_used` property into the top_metric parameter. If the metric type is not specified, by

default the API returns a list of directories with the greatest number of average read operations per second. The current maximum number of directories returned by the API for a metric type is 25.

#### Failure to return list of directories with most I/O activity

The API can sometimes fail to return the list of directories with the most I/O activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume has not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of directories with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return list of largest directories

The API can sometimes fail to return the list of largest directories, due to the following reasons:

– The volume does not have file system analytics enabled.

– The volume’s file system analytics database version doesn’t support this report.

#### Failure to return pathnames

The API can sometimes fail to obtain filesystem pathnames for certain directories, either due to internal transient errors or if those directories have been recently deleted.<br> In such cases, instead of the pathname, the API will return "{<volume_instance_uuid>.<fileid>}" for that directory.</fileid></volume_instance_uuid>

#### Examples

### Retrieving a list of the directories with the greatest average number of read operations per second

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=iops.read"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1",
```

```
    },
    "iops": {
      "read": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1505
      }
    },
    "path": "/dir1/dir2",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "vol1",
    },
    "iops": {
      "read": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "path": "/dir3/dir4",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
```

```
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir3%2Fdir4"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir3%2Fdir4?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
      "name": "vol1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864"
        }
      }
    },
    "iops": {
      "read": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "path": "/dir12",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
```

```
2df8e52ef864/files/dir12"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/directories?top_metric=iops.read"
  }
}
}
```

**Retrieving a list of the directories with the most traffic with failure in obtaining the pathnames for the directories:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=throughput.write"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 24,
      "error": {
        "lower_bound": 24,
        "upper_bound": 29
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:1232}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
```

```
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    }
  },
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 12,
      "error": {
        "lower_bound": 12,
        "upper_bound": 22
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:6754}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    }
  },
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 8,
      "error": {
        "lower_bound": 8,
        "upper_bound": 10
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:8654}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
```

```
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/top-metrics/directories?top_metric=throughput.write"
  }
}
}
```

#### Retrieving a list of the largest directories

The following example shows how to retrieve a list of the largest directories based on the non-recursive bytes used by the contents of a directory.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=non_recursive_bytes_used"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 345,
    "path": "/dir1/dir2",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
```

```
2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 345,
    "path": "/dir3/dir4",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir3%2Fdir4"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir3%2Fdir4?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 345,
    "path": "/dir12",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
```

```
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
  }
}
}
```


**Retrieving a list of the largest directories where incomplete data is reported**

The following example retrieves a list of the largest directories when partial data is returned.


```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=non_recursive_bytes_used"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 1022,
    "path": "/dir1/dir2",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
```

```
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 261,
    "path": "/dir12",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir12?return_metadata=true"
      }
    }
  }
],
"num_records": 2,
```

```
"incomplete_response_reason": {
  "message": "Partial data has been returned for this metric report.
Reason: Data collection for the large directory report is in progress.",
  "code": "111411234",
  "arguments": [
    {
      "message": "Data collection for the large directory report is in
progress."
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
  }
}
}
```


#### Retrieving a list of the largest directories when the file system analytics database version doesn’t support this report

The following example shows the behavior of the API when the file system analytics database version doesn’t support the large directory report.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=non_recursive_bytes_used"

# The Response:
{
"records": [
],
"num_records": 0,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
  }
},
"error": {
  "message": "The large directory report for volume \"FV\" in Vserver
\"vs0\" is not available because the file system analytics database
version doesn't support this report. Use the \"volume analytics off\"
command to disable analytics on the volume, then use the \"volume
analytics on\" command to re-enable analytics.",
  "code": "124519410"
}
}
```

**Example showing the behavior of the API when there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/directories?top_metric=throughput.write"
  }
}
}
```

**Retrieve directories with the greatest value performance metric or capacity metric**

GET `/storage/volumes/{volume.uuid}/top-metrics/directories`

**Introduced In:** 9.10

Retrieves a list of directories with the greatest value performance metric or capacity metric.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| top_metric | string | query | False | Type of performance metric or capacity metric.<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write", "non_recursive_bytes_used"] |
| max_records_per_volume | integer | query | False | Max records per volume. |
| non_recursive_bytes_used | integer | query | False | Filter by non_recursive_bytes_used<br>• Introduced in: 9.12 |
| path | string | query | False | Filter by path |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.write | integer | query | False | Filter by throughput.write |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| volume.name | string | query | False | Filter by volume.name |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| incomplete_response_reason | incomplete_response_reason | Indicates that the metric report provides incomplete data. |
| notice | notice | Optional field that indicates why no records are returned by the volume activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_directory] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: Data collection for the large directory report is in progress."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 10,
        "write": 5
      },
      "non_recursive_bytes_used": 300,
```

```
      "path": "/dir_abc/dir_123/dir_20",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 3,
        "write": 20
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |

| Error Code | Description |
| --- | --- |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518415 | Failed to get the volume activity tracking report on volume volume.name in SVM svm.name. |
| 124518416 | Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name. |
| 124518417 | Volume wildcard queries are not supported for activity tracking reports. |
| 124518418 | The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic. Refer to the REST API documentation for more information. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 124518423 | Volume activity tracking is not supported on the volume. |
| 124519410 | The large directory report for volume volume.name in SVM svm.name is not available because the file system analytics database version doesn't support this report. |
| 124519411 | Volume activity tracking is not enabled on the volume. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

incomplete_response_reason

Indicates that the metric report provides incomplete data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

notice

Optional field that indicates why no records are returned by the volume activity tracking REST API.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

_links

| Name | Type | Description |
|---|---|---|
| metadata | href | |
| self | href | |

top_metric_value_error_bounds

| Name | Type | Description |
|---|---|---|
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_directory

Information about a directory's IO metrics.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| iops | iops | |
| non_recursive_bytes_used | integer | Non-recursive bytes used by the contents of a directory. |
| path | string | Path of the directory. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |

| Name | Type | Description |
| --- | --- | --- |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View I/O activity for files

### View I/O activity for files

#### Overview

You can use this API to retrieve a list of files with the most I/O activity for a specified volume, within the past several seconds. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` or `throughput` metrics. This API also supports only returning one I/O activity type per request.

#### Approximate accounting and error bars

When too many files have recent activity, some files might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the files that would have otherwise been included might not be present in the final list returned by the API.

#### Failure to return list of files with most I/O activity

The API can sometimes fail to return the list of files with the most I/O activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume has not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of files with the most I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return pathnames

The API can sometimes fail to obtain the filesystem pathnames for certain files, either due to internal transient errors or if those files have been recently deleted.<br> In such cases, instead of the pathname, the API will return "{<volume_instance_uuid>.<fileid>}" for that file.</fileid></volume_instance_uuid>

#### Retrieve a list of the files with the most I/O activity

For a report on the files with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of the files with the greatest number of the average read operations per second. The current maximum number of files returned by the API for an I/O activity type is 25.

– GET /api/storage/volumes/{volume.uuid}/top-metrics/files

#### Examples

**Retrieving a list of the files with the greatest average number of write bytes received per second:**

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/files

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/files?top_metric=throughput.write"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 24,
      "error": {
        "lower_bound": 24,
        "upper_bound": 29
      }
    },
    "path": "/d5/f5",
```

```
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    },
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/files/d5%2Ff5?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 12,
      "error": {
        "lower_bound": 12,
        "upper_bound": 22
      }
    },
    "path": "/d6/f6",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    },
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/files/d6%2Ff6?return_metadata=true"
      }
    }
  },
  {
    "volume": {
```

```
      "name": "fv"
    },
    "throughput": {
      "write": 8,
      "error": {
        "lower_bound": 8,
        "upper_bound": 10
      }
    },
    "path": "/d3/f3",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    },
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/files/d3%2Ff3?return_metadata=true"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/top-metrics/files?top_metric=throughput.write"
  }
}
}
```


**Retrieving a list of the files with the most traffic with failure in obtaining the pathnames for the files:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/files?top_metric=throughput.write"

# The Response:
{
"records": [
```

```
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 24,
      "error": {
        "lower_bound": 24,
        "upper_bound": 29
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:1232}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    }
  },
  {
    "volume": {
      "name": "fv"
    },
    "throughput": {
      "write": 12,
      "error": {
        "lower_bound": 12,
        "upper_bound": 22
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:6754}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    }
  },
  {
    "volume": {
```

```
      "name": "fv"
    },
    "throughput": {
      "write": 8,
      "error": {
        "lower_bound": 8,
        "upper_bound": 10
      }
    },
    "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:8654}",
    "svm": {
      "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
      "name": "vs0",
      "_links": {
        "self": {
          "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
        }
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-
005056ac0f77/top-metrics/files?top_metric=throughput.write"
  }
}
}
```


**Example showing the behavior of the API when there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/files?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/files?top_metric=throughput.write"
  }
}
}
```

**Retrieve files with the most I/O activity**

GET `/storage/volumes/{volume.uuid}/top-metrics/files`

**Introduced In:** 9.10

Retrieves a list of files with the most I/O activity.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| max_records_per_volume | integer | query | False | Max records per volume. |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| volume.name | string | query | False | Filter by volume.name |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| path | string | query | False | Filter by path |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| incomplete_response_reason | incomplete_response_reason | Indicates that the metric report provides incomplete data. |
| notice | notice | Optional field that indicates why no records are returned by the volume activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_file] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: The activity tracking report for this volume is not available
because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 5,
        "write": 4
      },
```

```
      "path": "/dir_abc/dir_123/file_1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 2,
        "write": 20
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |

| Error Code | Description |
| --- | --- |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518415 | Failed to get the volume activity tracking report on volume volume.name in SVM svm.name. |
| 124518416 | Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name. |
| 124518417 | Volume wildcard queries are not supported for activity tracking reports. |
| 124518418 | The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic. Refer to the REST API documentation for more information. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 124518423 | Volume activity tracking is not supported on the volume. |
| 124519411 | Volume activity tracking is not enabled on the volume. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

incomplete_response_reason

Indicates that the metric report provides incomplete data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

notice

Optional field that indicates why no records are returned by the volume activity tracking REST API.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

_links

| Name | Type | Description |
|---|---|---|
| metadata | href | |
| self | href | |

top_metric_value_error_bounds

| Name | Type | Description |
|---|---|---|
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

volume

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_file

Information about a file's IO activity.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| iops | iops | |
| path | string | Path of the file. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| volume | volume | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View I/O activity for users

### View I/O activity for users

#### Overview

You can use this API to retrieve a list of users with the most I/O activity for a specified volume, within the past several seconds. Use the `top_metric` parameter to specify which type of I/O activity to filter for. This API is used to provide insight into I/O activity and supports ordering by I/O activity types, namely `iops` or `throughput` metrics. This API also supports only returning one I/O activity type per request.

#### Approximate accounting and error bars

When too many users have recent activity, some users might be dropped from the list. In this situation, the spread of values in the `error` field increases, indicating that there are larger error bars on the value for `iops` or `throughput`. As the list becomes increasingly more approximate due to dropped entries, some of the users that would have otherwise been included might not be present in the final list returned by the API.

#### Failure to return list of users with most I/O activity

The API can sometimes fail to return the list of users with the most I/O activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume has not had any recent NFS/CIFS client traffic.

– The NFS/CIFS client operations are being served by the client-side filesystem cache.

– The NFS/CIFS client operations are being buffered by the client operating system.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of users with the most

I/O activity.

– NFSv4 client read operations using Multi-Processor I/O (MPIO) are not tracked.

#### Failure to return the usernames

The API can sometimes fail to obtain the usernames for the list of userid entries, due to internal transient errors.<br> In such cases, instead of the username, the API will return "{<user-id>}" for the user entry.</user-id>

#### Retrieve a list of the users with the most I/O activity

For a report on the users with the most I/O activity returned in descending order, specify the I/O activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the I/O activity type is not specified, by default the API returns a list of the users with the greatest number of the average read operations per second. The current maximum number of users returned by the API for an I/O activity type is 25.

– GET /api/storage/volumes/{volume.uuid}/top-metrics/users

#### Examples

### Retrieving a list of the users with the greatest average number of read bytes received per second:

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/users

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/users?top_metric=throughput.read"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "throughput": {
      "read": 1495,
      "error": {
        "lower_bound": 1495,
        "upper_bound": 1502
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
    "user_name": "John",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
```

```
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "throughput": {
      "read": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1025
      }
    },
    "user_id": "1988",
    "user_name": "Ryan",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "throughput": {
      "read": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 348
      }
    },
    "user_id": "S-1-5-21-256008430-3394229847-3930036330-1003",
    "user_name": "Julie",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
```

```
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/users?top_metric=throughput.read"
  }
}
}
```

**Example showing the behavior of the API when there is no read/write traffic:**

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/users?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/users?top_metric=throughput.write"
  }
}
}
```

### Retrieve users with the most I/O activity

GET `/storage/volumes/{volume.uuid}/top-metrics/users`

**Introduced In:** 9.10

Retrieves a list of users with the most I/O activity.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| volume.uuid | string | path | True | Volume UUID |
| top_metric | string | query | False | I/O activity type<br>• Default value: 1<br>• enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"] |
| throughput.write | integer | query | False | Filter by throughput.write |
| throughput.read | integer | query | False | Filter by throughput.read |
| throughput.error.upper_bound | integer | query | False | Filter by throughput.error.upper_bound |
| throughput.error.lower_bound | integer | query | False | Filter by throughput.error.lower_bound |
| svm.name | string | query | False | Filter by svm.name |
| svm.uuid | string | query | False | Filter by svm.uuid |
| volume.name | string | query | False | Filter by volume.name |
| iops.error.upper_bound | integer | query | False | Filter by iops.error.upper_bound |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| iops.error.lower_bound | integer | query | False | Filter by iops.error.lower_bound |
| iops.read | integer | query | False | Filter by iops.read |
| iops.write | integer | query | False | Filter by iops.write |
| user_name | string | query | False | Filter by user_name |
| user_id | string | query | False | Filter by user_id |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| incomplete_response_reason | incomplete_response_reason | Indicates that the metric report provides incomplete data. |
| notice | notice | Optional field that indicates why no records are returned by the volume activity tracking REST API. |
| num_records | integer | Number of records. |
| partial_response_reason | partial_response_reason | Indicates that the metric report provides partial data. |
| records | array[top_metrics_user] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: The activity tracking report for this volume is not available
because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "partial_response_reason": {
    "code": "124518424",
    "message": "The top metrics report contains partial data for read
operations because NFSv4 reads using Multi-Processor I/O (MPIO) are not
tracked."
  },
  "records": [
    {
      "iops": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 4,
        "write": 8
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
      },
      "throughput": {
        "error": {
          "lower_bound": 34,
          "upper_bound": 54
        },
        "read": 10,
        "write": 7
      },
      "user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "user_name": "James",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 124518405 | Volume activity tracking is not supported on volumes that contain LUNs. |
| 124518407 | Volume activity tracking is not supported on FlexCache volumes. |
| 124518408 | Volume activity tracking is not supported on audit staging volumes. |
| 124518409 | Volume activity tracking is not supported on object store server volumes. |
| 124518410 | Volume activity tracking is not supported on SnapMirror destination volumes. |
| 124518411 | Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents. |

| Error Code | Description |
| --- | --- |
| 124518412 | Volume activity tracking is not supported on SnapLock volumes. |
| 124518414 | Volume activity tracking is not supported on volumes that contain NVMe namespaces. |
| 124518415 | Failed to get the volume activity tracking report on volume volume.name in SVM svm.name. |
| 124518416 | Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name. |
| 124518417 | Volume wildcard queries are not supported for activity tracking reports. |
| 124518418 | The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic. Refer to the REST API documentation for more information. |
| 124518422 | Volume activity tracking is not supported on All SAN Array clusters. |
| 124518423 | Volume activity tracking is not supported on the volume. |
| 124519411 | Volume activity tracking is not enabled on the volume. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

incomplete_response_reason

Indicates that the metric report provides incomplete data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

notice

Optional field that indicates why no records are returned by the volume activity tracking REST API.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why no records are returned. |
| message | string | Details why no records are returned. |

partial_response_reason

Indicates that the metric report provides partial data.

| Name | Type | Description |
|---|---|---|
| code | string | Warning code indicating why partial data was reported. |
| message | string | A message describing the reason for partial data. |

top_metric_value_error_bounds

| Name | Type | Description |
|---|---|---|
| lower_bound | integer | Lower bound of the nominal value of a metric. |
| upper_bound | integer | Upper bound of the nominal value of a metric. |

iops

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |
| read | integer | Average number of read operations per second. |
| write | integer | Average number of write operations per second. |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

throughput

| Name | Type | Description |
|---|---|---|
| error | top_metric_value_error_bounds | |

| Name | Type | Description |
| --- | --- | --- |
| read | integer | Average number of read bytes received per second. |
| write | integer | Average number of write bytes received per second. |

volume

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | The name of the volume. This field cannot be specified in a PATCH method. |
| uuid | string | Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.<br>• example: 028baa66-41bd-11e9-81d5-00a0986138f7<br>• Introduced in: 9.6<br>• x-nullable: true |

top_metrics_user

Information about a user's IO activity.

| Name | Type | Description |
| --- | --- | --- |
| iops | iops | |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| throughput | throughput | |
| user_id | string | User ID of the user. |
| user_name | string | Name of the user. |
| volume | volume | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Support

## ONTAP REST API Support endpoints

### Overview

You can use ONTAP support APIs to manage configuration backups, autosupport settings, and event handling.

---

### Configuration backups

Configuration backups are copies of node and cluster settings saved to an external server.
Single-node clusters must have configuration backups as protection against corruption of the configuration database. Multi-node clusters back up the configuration automatically between the nodes in the cluster.

---

### Event Management System (EMS)

The Event Management System (EMS) collects and processes events, and sends notification of the events through various reporting mechanisms.

### Examples

The following example configures EMS to send a support email when a WAFL event is observed with a severity of error, alert, or emergency.

**Configuring the system-wide email parameters**

```
# The API:
PATCH /support/ems

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/ems" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_config_patch.txt"
test_ems_config_patch.txt(body):
{
"mail_from": "administrator@mycompany.com",
"mail_server": "smtp@mycompany.com"
}

# The response:
200 OK
```

**Configuring a filter with an enclosed rule**

```
# The API:
POST /support/ems/filters

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/filters" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_filters_post.txt"
test_ems_filters_post.txt(body):
{
"name": "error-wafl",
"rules": [
  {
    "index": 1,
    "type": "include",
    "message_criteria": {
      "name_pattern": "wafl.*",
      "severities": "emergency,alert,error"
    }
  }
]
}

# The response:
201 Created
```

#### Setting up an email destination

```
# The API:
POST /support/ems/destinations

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/destinations" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_dest_post.txt"
test_ems_dest_post.txt(body):
{
"name": "Technician_Email",
"type": "email",
"destination": "technician@mycompany.com",
"filters": [
  { "name" : "error-wafl" }
]
}

# The response:
201 Created
```

## Manage an automatic update feature

### Manage an automatic update feature

#### Overview

The Automatic Update feature helps keep the ONTAP cluster current with the latest software component updates by automatically downloading and applying them to the cluster.
By enabling this feature, you agree to the following terms:

---

### AUTOMATIC UPDATE TERMS

*These Automatic Update Terms ("Terms") set forth the terms and conditions between NetApp, Inc., NetApp B.V., or any of their affiliates ("NetApp") and End User Customer ("Customer") in connection with the feature enabling Customer to receive software patches, upgrades, and updates to NetApp Software automatically ("Automatic Update"). By agreeing to and accepting Automatic Updates, Customer agrees to be bound by these Terms, as well as NetApp's End User License Agreement and Support Terms available at https://www.netapp.com/how-to-buy/sales-terms-and-conditions/.*

*By enabling the Automatic Update feature, Customer agrees to receive Automatic Updates that NetApp may provide from time to time, without any additional notice, and NetApp will not be liable for any damages, loss of data or loss of functionalities arising from provision of Automatic Updates. Customer may revoke acceptance of these Terms and disable the receipt of Automatic Updates by setting the feature configuration to "Disabled" in ONTAP.*

*Timely installation of updates to security features is required for such features to function properly, including*

*ransomware detection.*

---

Important note:
When the automatic update feature is disabled

- No new updates are shown to the user
- All automatic updates currently scheduled will have their schedules cancelled.
- All automatic updates currently waiting for user confirmation cannot be started until the feature is re-enabled.

---

#### Examples

### Retrieving the current status of the automatic update feature

The following example shows how to retrieve the current status of the automatic update feature:

```
# The API:
GET /support/auto-update

# The call:
curl -X GET "https://<mgmt-ip>/api/support/auto-update"

# The response:
{
"enabled": true,
"eula": {
  "accepted": true,
  "accepted_timestamp": "2020-12-01T21:24:44-04:00",
  "user_id_accepted": "admin",
  "accepted_ip_address": "192.168.1.125"
},
"_links": {
  "self": {
    "href": "/api/support/auto-update"
  }
}
}
```

---

### Updating the status of the automatic update feature

The following example shows how to update the status of the automatic update feature:

```
# The API:
PATCH /support/auto-update

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/auto-update" -H "accept:
application/hal+json" -d '{ "enabled" : "true" }'

# The response:
200 OK
```

**Retrieve the current status of an automatic update feature**

GET `/support/auto-update`

**Introduced In:** 9.10

Retrieves the current status of the automatic update feature and the End User License Agreement (EULA).

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Flag indicating feature state. |
| eula | eula | |

### Example response

```
{
  "_links": {
    "self": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "enabled": 1,
  "eula": {
    "accepted": 1,
    "accepted_ip_address": "192.168.1.125",
    "accepted_timestamp": "2020-12-01 09:12:23 -0400",
    "user_id_accepted": "admin"
  }
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

eula

| Name | Type | Description |
|---|---|---|
| accepted | boolean | Flag indicating the End User License Agreement (EULA) acceptance. When the feature is enabled, it is assumed that the EULA is accepted. |
| accepted_ip_address | string | IP Address from where the EULA was accepted. |
| accepted_timestamp | string | Date and time when the EULA was accepted. |
| user_id_accepted | string | User ID that provided the EULA acceptance. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the current enabled status of an automatic update feature**

PATCH `/support/auto-update`

**Introduced In:** 9.10

Updates the current enabled status of the automatic update feature and accepts the EULA.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Set the force flag to true to enable the automatic update feature regardless of how AutoSupport is configured. Without this flag set to true, an attempt to enable the automatic update feature fails unless AutoSupport is enabled, its transport is HTTPS, and the AutoSupport OnDemand feature is enabled.<br>• Introduced in: 9.16<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Flag indicating feature state. |
| eula | eula | |

#### Example request

```
{
  "enabled": 1,
  "eula": {
    "accepted": 1,
    "accepted_ip_address": "192.168.1.125",
    "accepted_timestamp": "2020-12-01 09:12:23 -0400",
    "user_id_accepted": "admin"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262179 | Unexpected argument. |
| 131072216 | Automatic update requires AutoSupport to be enabled. |
| 131072217 | Automatic update requires AutoSupport OnDemand to be enabled. |
| 131072218 | Automatic update requires AutoSupport to use the HTTPS transport. |
| 131072219 | The provided parameter requires an effective cluster version of ONTAP 9.16.1 or later. |

Also see the table of common errors in the Response body overview section of this documentation.

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self

_links

eula

| Name | Type | Description |
|---|---|---|
| accepted | boolean | Flag indicating the End User License Agreement (EULA) acceptance. When the feature is enabled, it is assumed that the EULA is accepted. |
| accepted_ip_address | string | IP Address from where the EULA was accepted. |
| accepted_timestamp | string | Date and time when the EULA was accepted. |
| user_id_accepted | string | User ID that provided the EULA acceptance. |

auto_update_info

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| enabled | boolean | Flag indicating feature state. |
| eula | eula | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve automatic update configuration

### Retrieve automatic update configuration

#### Overview

Use this API to query and retrieve all the configurations for automatic updates.

This API only supports GET calls.

---

#### Examples

### Retrieving the configuration settings for automatic updates

The following example shows how to retrieve settings for all automatic updates:

```
# The API:
GET /support/auto-update/configurations/

# The call:
curl -X GET "https://<mgmt-ip>/api/support/auto-
update/configurations?fields=*"

# The response:
{
"records": [
  {
    "uuid": "440ae2e4-fd8f-4225-9bee-94e2da3f8d9d",
    "category": "firmware",
    "description": {
      "message": "SP/BMC Firmware",
      "code": "131072402"
    },
    "action": "confirm",
  }
],
"num_records": 1
}
```

**Retrieve the configuration for an automatic update**

GET `/support/auto-update/configurations`

**Introduced In:** 9.10

Retrieves the configuration for automatic updates.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| category | string | query | False | Filter by category |
| action | string | query | False | Filter by action |
| uuid | string | query | False | Filter by uuid |
| description.message | string | query | False | Filter by description.message |
| description.code | string | query | False | Filter by description.code |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[auto_update_configuration] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "action": "confirm",
      "category": "disk_fw",
      "description": {
        "code": "string",
        "message": "string"
      },
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff47"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

description

Description of the configuration row.
ONTAP Message Codes

| Code | Description |
|---|---|
| 131072401 | Storage Firmware |
| 131072402 | SP/BMC Firmware |
| 131072403 | System Files |

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

auto_update_configuration

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| action | string | The action to be taken by the alert source as specified by the user. |
| category | string | Category for the configuration row. |

| Name | Type | Description |
| --- | --- | --- |
| description | description | Description of the configuration row.<br><br>ONTAP Message Codes<br><br>Code — Description<br><br>131072401 — Storage Firmware<br><br>131072402 — SP/BMC Firmware<br><br>131072403 — System Files |
| uuid | string | Unique identifier for the configuration row. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage automatic update configuration

### Manage automatic update configuration

#### Overview

Use this API to query and retrieve a specific automatic package update configuration.

This API supports GET and PATCH calls. PATCH enables the `action` field to be updated for the specified configuration.

**Examples**

### Retrieving settings for a specific automatic update

The following example shows how to retrieve configuration settings for an automatic update category:

```
# The API:
/support/auto-update/configurations/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/auto-
update/configurations/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d"

# The response:
{
"uuid": "440ae2e4-fd8f-4225-9bee-94e2da3f8d9d",
"category": "firmware",
"description": {
  "message": "SP/BMC Firmware",
  "code": "131072402"
},
"action": "confirm",
}
```

### Updating the settings for a specific automatic update

The following example shows how to modify configuration settings for an automatic update:

```
# The API:
/support/auto-update/configurations/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/auto-
update/configurations/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d" -H "accept:
application/hal+json" -d '{"action":"confirm"}'

# The response:
200 OK
```

### Retrieve a specific configuration for an automatic update

GET `/support/auto-update/configurations/{uuid}`

**Introduced In:** 9.10

Retrieves the configuration for a specified automatic update.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Unique identifier for configuration record. |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| action | string | The action to be taken by the alert source as specified by the user. |
| category | string | Category for the configuration row. |
| description | description | Description of the configuration row.<br><br>ONTAP Message Codes<br><br>Code — Description<br><br>131072401 — Storage Firmware<br><br>131072402 — SP/BMC Firmware<br><br>131072403 — System Files |
| uuid | string | Unique identifier for the configuration row. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "action": "confirm",
  "category": "disk_fw",
  "description": {
    "code": "string",
    "message": "string"
  },
  "uuid": "572361f3-e769-439d-9c04-2ba48a08ff47"
}
```

### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

description

Description of the configuration row.
ONTAP Message Codes

| Code | Description |
|---|---|
| 131072401 | Storage Firmware |
| 131072402 | SP/BMC Firmware |
| 131072403 | System Files |

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Modify the configuration for an automatic update

PATCH `/support/auto-update/configurations/{uuid}`

**Introduced In:** 9.10

Updates the configuration for a specified automatic update.

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| uuid | string | path | True | Unique identifier for configuration record. |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| action | string | The action to be taken by the alert source as specified by the user. |
| category | string | Category for the configuration row. |
| uuid | string | Unique identifier for the configuration row. |

### Example request

```
{
  "action": "confirm",
  "category": "disk_fw",
  "uuid": "572361f3-e769-439d-9c04-2ba48a08ff47"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262179 | Unexpected argument. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

description

Description of the configuration row.
ONTAP Message Codes

| Code | Description |
|---|---|
| 131072401 | Storage Firmware |
| 131072402 | SP/BMC Firmware |
| 131072403 | System Files |

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

auto_update_configuration

| Name | Type | Description |
|---|---|---|
| action | string | The action to be taken by the alert source as specified by the user. |
| category | string | Category for the configuration row. |
| uuid | string | Unique identifier for the configuration row. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve status for all automatic updates

### Retrieve status for all automatic updates

#### Overview

Use this API to retrieve the status for all the automatic package updates.

This API only supports GET calls.

---

#### Examples

### Retrieving the status of all package updates

The following example shows how to retrieve the status of all automatic updates:

```
# The API:
GET /support/auto-update/updates/

# The call:
curl -X GET "https://<mgmt-ip>/api/support/auto-update/updates?fields=*"

# The response:
{
"records": [
  {
    "uuid" : "440ae2e4-fd8f-4225-9bee-94e2da3f8d9d",
    "package_id" : "572361f3-e769-439d-9c04-2ba48a08ff47",
    "content_type" : "disk_fw",
    "content_category" : "firmware",
    "description" : "disk_fw version 3.0",
    "state" : "downloading",
    "creation_time" : "2020-12-01T09:12:03Z",
```

```
    "expiry_time" : "2021-06-01T09:12:03Z",
    "last_state_change_time" : "2020-12-01T09:12:23Z",
    "start_time": "2020-12-01T09:12:23Z",
    "percent_complete": 25,
    "remaining_time": "PT1M30S",
    "_links": {
      "self": {
        "href": "/support/auto-update/updates/440ae2e4-fd8f-4225-9bee-
94e2da3f8d9d"
      }
    }
  },
  {
    "uuid" : "440ae2e4-fd8f-4225-9bee-94e2da3f9d8d",
    "package_id" : "e2e12753-8501-11e8-a7c1-00505695c997",
    "content_type" : "disk_fw",
    "content_category" : "firmware",
    "description" : "disk_fw version 3.1",
    "state" : "pending_confirmation",
    "creation_time" : "2020-12-01T09:12:03Z",
    "expiry_time" : "2021-06-01T09:12:03Z",
    "last_state_change_time" : "2020-12-01T09:12:23Z",
    "start_time": "2020-12-01T09:12:23Z",
    "_links": {
      "self": {
        "href": "/support/auto-update/updates/440ae2e4-fd8f-4225-9bee-
94e2da3f8d9d"
      }
    }
  },
  {
    "uuid" : "440ae2e4-fd8f-4225-9aeb-94e2da3f8ef9",
    "package_id" : "e2e12753-8501-11e8-a7c1-00505695c997",
    "content_type" : "shelf_fw",
    "content_category" : "Firmware",
    "description" : "shelf_fw version 4.0",
    "state" : "applying",
    "creation_time" : "2020-12-01T09:11:53Z",
    "expiry_time" : "2021-06-01T09:12:03Z",
    "last_state_change_time" : "2020-12-01T09:12:13Z",
    "start_time": "2020-12-01T09:12:23Z",
    "percent_complete": 85,
    "remaining_time": "PT30S",
    "_links": {
      "self": {
        "href": "/support/auto-update/updates/440ae2e4-fd8f-4225-9aeb-
```

```
94e2da3f8ef9"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/support/auto-update/updates?fields=*"
  }
}
}
```


**Retrieve status of all automatic updates**

GET `/support/auto-update/updates`

**Introduced In:** 9.10

Retrieves the status of all updates.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| remaining_time | string | query | False | Filter by remaining_time |
| scheduled_time | string | query | False | Filter by scheduled_time |
| uuid | string | query | False | Filter by uuid |
| description | string | query | False | Filter by description |
| state | string | query | False | Filter by state |
| last_state_change_time | string | query | False | Filter by last_state_change_time |
| start_time | string | query | False | Filter by start_time |
| creation_time | string | query | False | Filter by creation_time |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| content_category | string | query | False | Filter by content_category |
| package_id | string | query | False | Filter by package_id |
| end_time | string | query | False | Filter by end_time |
| expiry_time | string | query | False | Filter by expiry_time |
| percent_complete | integer | query | False | Filter by percent_complete |
| content_type | string | query | False | Filter by content_type |
| status.arguments.message | string | query | False | Filter by status.arguments.message |
| status.arguments.code | string | query | False | Filter by status.arguments.code |
| status.code | string | query | False | Filter by status.code |
| status.message | string | query | False | Filter by status.message |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[auto_update_status] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "action": "schedule_now",
      "content_category": "Firmware",
      "content_type": "disk_fw",
      "creation_time": "2020-12-01 09:12:23 +0000",
      "description": "disk_fw version 3.0",
      "end_time": "2020-12-01 09:12:23 +0000",
      "expiry_time": "2021-06-01 09:12:23 +0000",
      "last_state_change_time": "2020-12-01 09:12:23 +0000",
      "package_id": "572361f3-e769-439d-9c04-2ba48a08ff47",
      "percent_complete": 85,
      "remaining_time": "PT1H45M13S",
      "schedule_time": "2020-12-20 21:00:00 +0000",
      "scheduled_time": "2020-12-05 09:12:23 +0000",
      "start_time": "2020-12-01 09:12:23 +0000",
      "state": "pending_confirmation",
      "status": {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      },
```

```
      "uuid": "440ae2e4-fd8f-4225-9bee-94e2da3f9d8d"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

status

Status of the update.
ONTAP Message Codes

| Code | Description |
|---|---|
| 8650878 | Get-url request to AutoSupport OnDemand Server failed. |
| 131072105 | Package installation successful. |
| 131072106 | Package installation failed. |
| 131072107 | Package installation is in progress. |
| 131072108 | An internal error occurred. |
| 131072109 | No errors found. |

| Code | Description |
| --- | --- |
| 131072110 | Failed to extract package. |
| 131072111 | Failed to create entry in status table. |
| 131072209 | Failed to download package. |
| 131072210 | Failed to obtain download links. Will be retried. |
| 131072211 | Installation cannot start because the following nodes are not up. Will be retried. |
| 131072212 | Automatic update cannot start. Will be retried. |
| 131072213 | Installation complete. |
| 131072214 | Installation in progress. |
| 131072215 | Scheduled to start. |
| 131072301 | Unable to create directory. |
| 131072302 | Update not required. It is already on the target firmware version. |
| 131072303 | Unable to move installable to installation location. |
| 131072304 | Successfully updated. All nodes already updated and skipped. |
| 131072306 | Firmware distribution starting. |
| 131072307 | Firmware distribution restarting. |
| 131072308 | Firmware distribution failed. |
| 131072309 | Firmware installation successful on node. |
| 131072310 | Firmware installable is corrupt on node. |
| 131072311 | Could not apply firmware update on nodes. |
| 131072312 | Firmware update operation failed. |
| 131072313 | Install failed to initiate on node. |
| 131072314 | Install initiated on node. |
| 131072315 | Install failed on node. |
| 131072316 | Update in progress on node. |
| 131072317 | Unexpected state occurred during firmware installation on node. |
| 131072318 | Successfully updated SP/BMC firmware. |
| 131072319 | Backup SP image is running. Firmware installation failed prematurely on node. |
| 131072320 | SP rebooted with an unexpected firmware version on node. |
| 131072321 | Waiting for shelf firmware to update. |

| Code | Description |
| --- | --- |
| 131072322 | Update complete. |
| 131072323 | Shelf update timed out. |
| 131072324 | Successfully updated shelves on nodes. |
| 131072325 | Waiting for disk firmware to update. |
| 131072326 | Disk update completed. |
| 131072327 | Disk update timed out. |
| 131072328 | Successfully updated disks on nodes. |
| 131072329 | Waiting for DQP firmware to update. |
| 131072330 | Successfully updated DQP on nodes. |
| 131072331 | DQP update timed out. |
| 131072332 | Failed to copy firmware. |

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

auto_update_status

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| action | string | Action to be applied to the automatic update. |
| content_category | string | Category of the update |
| content_type | string | Image or package type. |
| creation_time | string | The date and time at which the update request was received. |
| description | string | Description of the update. |
| end_time | string | The date and time at which the update request processing ended. |

| Name | Type | Description |
| --- | --- | --- |
| expiry_time | string | The date and time at which the update request will expire. |
| last_state_change_time | string | The date and time at which the state of the update changed last. |
| package_id | string | Unique identifier provided by the back-end. |
| percent_complete | integer | Percentage of update completed |
| remaining_time | string | The time remaining for the update processing to complete in an ISO-8601 duration formatted string. |
| schedule_time | string | Date and time when an automatic update action is scheduled. This field is required when the action field is set to "schedule". |
| scheduled_time | string | The date and time at which the update request is currently scheduled for. |
| start_time | string | The date and time at which the update request processing started. |
| state | string | Current state of the update. |

| Name | Type | Description |
| --- | --- | --- |
| status | status | Status of the update.<br><br>ONTAP Message Codes<br><br>Code — Description<br><br>8650878 — Get-url request to AutoSupport OnDemand Server failed.<br><br>131072105 — Package installation successful.<br><br>131072106 — Package installation failed.<br><br>131072107 — Package installation is in progress.<br><br>131072108 — An internal error occurred.<br><br>131072109 — No errors found.<br><br>131072110 — Failed to extract package.<br><br>131072111 — Failed to create entry in status table.<br><br>131072209 — Failed to download package.<br><br>131072210 — Failed to obtain download links. Will be retried.<br><br>131072211 — Installation cannot start because the following nodes are not up. Will be retried.<br><br>131072212 — Automatic update cannot start. Will be retried.<br><br>131072213 — Installation complete.<br><br>131072214 — Installation in progress.<br><br>131072215 — Scheduled to start.<br><br>131072301 — Unable to create directory.<br><br>131072302 — Update not required. It is already on the target firmware version. |

| Name | Type | Description |
|---|---|---|
| uuid | string | Unique identifier for this update. |
| returned_error | | 131072304 — Successfully updated. All nodes already |

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage automatic updates

### Manage automatic updates

#### Overview

Use this API to retrieve the status for a specific automatic package update.

This API supports GET and PATCH calls. PATCH can be used to perform an action on an automatic update.

#### Examples

### Retrieving the status of an update

The following example shows how to retrieve the status of an automatic update:

131072310 — Firmware installable is corrupt on node.

131072311 — Could not apply firmware update on nodes.

131072312 — Firmware update operation failed.

131072313 — Install failed to initiate on node.

131072314 — Install initiated on node.

131072315 — Install failed on node.

131072316 — Update in progress on node.

131072317 — Unexpected state occurred during firmware installation on node.

131072318 — Successfully updated SP/BMC firmware.

131072319 — Backup SP image is running. Firmware installation failed prematurely on node.

131072320 — SP rebooted with an unexpected firmware version on node.

131072321 — Waiting for shelf firmware to update.

```
# The API:
GET /support/auto-update/updates/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d

# The call:
curl -X GET "https://<mgmt-ip>/api/support/auto-update/updates/440ae2e4-
fd8f-4225-9bee-94e2da3f8d9d"

# The response:
{
"uuid" : "440ae2e4-fd8f-4225-9bee-94e2da3f8d9d",
"package_id" : "572361f3-e769-439d-9c04-2ba48a08ff47",
"content_type" : "disk_fw",
"content_category" : "Firmware",
"description" : "disk_fw version 3.0",
"state" : "downloading",
"creation_time" : "2020-12-01T09:12:03Z",
"expiry_time" : "2021-06-01T09:12:03Z",
"last_state_change_time" : "2020-12-01T09:12:23Z",
"start_time" : "2020-12-01T09:12:23Z",
"percent_complete": 25,
"remaining_time": "PT1M30S",
"status": {
    "message": "Get-url request to AutoSupport OnDemand Server failed.
Error: Couldn't connect to server.",
    "code": "8650878"
},
"_links": {
  "self": {
    "href": "/support/auto-update/updates/440ae2e4-fd8f-4225-9bee-
94e2da3f8d9d"
  }
}
}
```

---

#### Updating the state of an automatic update

The following example shows how to trigger an automatic update that is waiting for user confirmation:

```
# The API:
PATCH /support/auto-update/updates/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/auto-update/updates/440ae2e4-
fd8f-4225-9bee-94e2da3f8d9d?action=schedule_now" -H "accept:
application/hal+json"

# The response:
200 OK
```

The following example shows how to dismiss an automatic update that is waiting for user confirmation:

```
# The API:
PATCH /support/auto-update/updates/440ae2e4-fd8f-4225-9bee-94e2da3f8d9d

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/auto-update/updates/440ae2e4-
fd8f-4225-9bee-94e2da3f8d9d?action=dismiss" -H "accept:
application/hal+json"

# Response
200 OK
```

### Retrieve the status of an update

GET `/support/auto-update/updates/{uuid}`

**Introduced In:** 9.10

Retrieves the status of an update.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Update identifier |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| action | string | Action to be applied to the automatic update. |
| content_category | string | Category of the update |
| content_type | string | Image or package type. |
| creation_time | string | The date and time at which the update request was received. |
| description | string | Description of the update. |
| end_time | string | The date and time at which the update request processing ended. |
| expiry_time | string | The date and time at which the update request will expire. |
| last_state_change_time | string | The date and time at which the state of the update changed last. |
| package_id | string | Unique identifier provided by the back-end. |
| percent_complete | integer | Percentage of update completed |
| remaining_time | string | The time remaining for the update processing to complete in an ISO-8601 duration formatted string. |
| schedule_time | string | Date and time when an automatic update action is scheduled.<br>This field is required when the action field is set to "schedule". |
| scheduled_time | string | The date and time at which the update request is currently scheduled for. |
| start_time | string | The date and time at which the update request processing started. |
| state | string | Current state of the update. |

| Name | Type | Description |
| --- | --- | --- |
| status | status | Status of the update.<br><br>ONTAP Message Codes<br><br>Code — Description<br><br>8650878 — Get-url request to AutoSupport OnDemand Server failed.<br><br>131072105 — Package installation successful.<br><br>131072106 — Package installation failed.<br><br>131072107 — Package installation is in progress.<br><br>131072108 — An internal error occurred.<br><br>131072109 — No errors found.<br><br>131072110 — Failed to extract package.<br><br>131072111 — Failed to create entry in status table.<br><br>131072209 — Failed to download package.<br><br>131072210 — Failed to obtain download links. Will be retried.<br><br>131072211 — Installation cannot start because the following nodes are not up. Will be retried.<br><br>131072212 — Automatic update cannot start. Will be retried.<br><br>131072213 — Installation complete.<br><br>131072214 — Installation in progress.<br><br>131072215 — Scheduled to start.<br><br>131072301 — Unable to create directory.<br><br>131072302 — Update not required. It is already on the target firmware version. |
| | | 131072303 — Unable to move |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | Unique identifier for this update. |

131072304 — Successfully updated. All nodes already updated and skipped.

**Example response**


```
{
  "_links": {
    "self": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "action": "schedule_now",
  "content_category": "Firmware",
  "content_type": "disk_fw",
  "creation_time": "2020-12-01 09:12:23 +0000",
  "description": "disk_fw version 3.0",
  "end_time": "2020-12-01 09:12:23 +0000",
  "expiry_time": "2021-06-01 09:12:23 +0000",
  "last_state_change_time": "2020-12-01 09:12:23 +0000",
  "package_id": "572361f3-e769-439d-9c04-2ba48a08ff47",
  "percent_complete": 85,
  "remaining_time": "PT1H45M13S",
  "schedule_time": "2020-12-20 21:00:00 +0000",
  "scheduled_time": "2020-12-05 09:12:23 +0000",
  "start_time": "2020-12-01 09:12:23 +0000",
  "state": "pending_confirmation",
  "status": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist"
  },
  "uuid": "440ae2e4-fd8f-4225-9bee-94e2da3f9d8d"
}
```


131072320 — SP rebooted with an unexpected firmware version on node.

131072321 — Waiting for shelf firmware to update.

131072322 — Update complete.

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

status

Status of the update.
ONTAP Message Codes

| Code | Description |
|---|---|
| 8650878 | Get-url request to AutoSupport OnDemand Server failed. |
| 131072105 | Package installation successful. |
| 131072106 | Package installation failed. |
| 131072107 | Package installation is in progress. |
| 131072108 | An internal error occurred. |
| 131072109 | No errors found. |
| 131072110 | Failed to extract package. |
| 131072111 | Failed to create entry in status table. |
| 131072209 | Failed to download package. |
| 131072210 | Failed to obtain download links. Will be retried. |

| Code | Description |
| --- | --- |
| 131072211 | Installation cannot start because the following nodes are not up. Will be retried. |
| 131072212 | Automatic update cannot start. Will be retried. |
| 131072213 | Installation complete. |
| 131072214 | Installation in progress. |
| 131072215 | Scheduled to start. |
| 131072301 | Unable to create directory. |
| 131072302 | Update not required. It is already on the target firmware version. |
| 131072303 | Unable to move installable to installation location. |
| 131072304 | Successfully updated. All nodes already updated and skipped. |
| 131072306 | Firmware distribution starting. |
| 131072307 | Firmware distribution restarting. |
| 131072308 | Firmware distribution failed. |
| 131072309 | Firmware installation successful on node. |
| 131072310 | Firmware installable is corrupt on node. |
| 131072311 | Could not apply firmware update on nodes. |
| 131072312 | Firmware update operation failed. |
| 131072313 | Install failed to initiate on node. |
| 131072314 | Install initiated on node. |
| 131072315 | Install failed on node. |
| 131072316 | Update in progress on node. |
| 131072317 | Unexpected state occurred during firmware installation on node. |
| 131072318 | Successfully updated SP/BMC firmware. |
| 131072319 | Backup SP image is running. Firmware installation failed prematurely on node. |
| 131072320 | SP rebooted with an unexpected firmware version on node. |
| 131072321 | Waiting for shelf firmware to update. |
| 131072322 | Update complete. |
| 131072323 | Shelf update timed out. |
| 131072324 | Successfully updated shelves on nodes. |
| 131072325 | Waiting for disk firmware to update. |

| Code | Description |
|---|---|
| 131072326 | Disk update completed. |
| 131072327 | Disk update timed out. |
| 131072328 | Successfully updated disks on nodes. |
| 131072329 | Waiting for DQP firmware to update. |
| 131072330 | Successfully updated DQP on nodes. |
| 131072331 | DQP update timed out. |
| 131072332 | Failed to copy firmware. |

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Perform an action on the update**

PATCH `/support/auto-update/updates/{uuid}`

**Introduced In:** 9.10

Perform an action on the update.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| uuid | string | path | True | Update identifier |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| action | string | query | True | Action to be performed on the update<br><br>• enum: ["cancel_schedule", "dismiss", "schedule", "schedule_now", "abort", "undismiss"] |
| schedule_time | string | query | False | Date and Time for which update is to be scheduled. This parameter is required only when action is "schedule". It should be in the ISO-8601 Date and Time format. Example- "2020-12-05T09:12:23Z".<br><br>• format: date-time |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 262248 | Unexpected parameter missing a delimiter. |
| 131072200 | Invalid Automatic Update ID. |
| 131072201 | No schedule to cancel for automatic update ID. |
| 131072202 | Cannot abort an automatic update that is not in progress. |
| 131072203 | Cannot dismiss a non-pending automatic update. |

| Error Code | Description |
| --- | --- |
| 131072204 | Cannot reactivate an automatic update that has not been dismissed. |
| 131072205 | Cannot schedule a non-pending automatic update to start. |
| 131072206 | Cannot schedule an automatic update past its expiration date. |
| 131072207 | Cannot schedule an automatic update in the past. |
| 131072208 | Cannot schedule automatic update while the feature is disabled. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage AutoSupport configuration

### Manage AutoSupport configuration

#### Overview

AutoSupport is the NetApp *call home* mechanism. AutoSupport sends configuration details, status details, and error reporting details to NetApp.

This endpoint supports both GET and PATCH calls. GET is used to retrieve AutoSupport configuration details for the cluster and PATCH is used to modify the AutoSupport configuration of the cluster. You can also use GET calls to check AutoSupport connectivity.

---

#### Examples

### Configuring 'to' addresses

The following example configures AutoSupport to send emails to 'to' addresses.

```
# The API:
PATCH /support/autosupport

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"to\": [
\"abc@netapp.com\", \"xyz@netapp.com\" ]}"

# The response:
200 OK
{}
```

---

### Configuring 'SMTP' transport

The following example configures AutoSupport to use 'SMTP' transport. The default transport is 'HTTPS'.

```
# The API:
PATCH /support/autosupport

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"transport\": \"smtp\"}"

# The response:
200 OK
{}
```

---

### Configuring 'start_tls' SMTP encryption

The following example configures AutoSupport to use 'start_tls' SMTP encryption. The default value for the smtp_encryption field is 'none'.

---

```
# The API:
PATCH /support/autosupport

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"smtp_encryption\": \"start_tls\"}"

# The response:
200 OK
{}
```

---

#### Retrieving the AutoSupport configuration

The following example retrieves AutoSupport configuration for the cluster.

```
# The API:
GET /support/autosupport

# The call:
curl -X GET "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json"

# The response:
200 OK
{
"enabled": true,
"mail_hosts": [
  "mailhost"
],
"from": "Postmaster",
"to": [
  "abc@netapp.com",
  "xyz@netapp.com"
],
"contact_support": true,
"transport": "smtp",
"proxy_url": "",
"is_minimal": false,
"smtp_encryption": "none",
"ondemand_enabled": true,
"_links": {
  "self": {
    "href": "/api/support/autosupport"
  }
}
}
```

#### Retrieving AutoSupport connectivity issues

The following example retrieves AutoSupport connectivity issues for the cluster. The `fields=issues` parameter must be specified, for the response to return connectivity issues. The `corrective_action` section might contain commands which needs to be executed on the ONTAP CLI.

Note that the connectivity check can take up to 10 seconds to complete.

```
# The API:
GET /support/autosupport

# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/support/autosupport?fields=issues" -H
"accept: application/hal+json"

# The response:
200 OK
{
"issues": [
  {
    "destination": "mailhost",
    "component": "mail_server",
    "node": {
      "name": "node3",
      "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
        }
      }
    },
    "issue": {
      "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'",
      "code": "53149746"
    },
    "corrective_action": {
      "message": "Check the hostname of the SMTP server",
      "code": "53149746"
    }
  },
  {
    "destination": "https://support.netapp.com/aods/asupmessage",
    "component": "ondemand_server",
    "node": {
      "name": "node3",
      "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
        }
      }
    },
    "issue": {
      "message": "AutoSupport OnDemand is disabled when \"-transport\" is
not set to \"https\".",
```

```
      "code": "53149740"
    },
    "corrective_action": {
      "message": "Run \"system node autosupport modify -transport https
-node <node name>\" to set \"-transport\" to \"https\".",
      "code": "53149740"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/autosupport"
  }
}
}
```


#### Retrieving AutoSupport configuration and connectivity issues

The following example retrieves AutoSupport configuration and connectivity issues on the cluster. Use `fields=*,issues` parameter to return both configuration and connectivity issues.


```
# The API:
GET /support/autosupport

# The call:
curl -X GET "https://<mgmt-ip>/api/support/autosupport?fields=*%2Cissues"
-H "accept: application/hal+json"

# The response:
200 OK
{
"enabled": true,
"mail_hosts": [
  "mailhost"
],
"from": "Postmaster",
"to": [
  "abc@netapp.com",
  "xyz@netapp.com"
],
"contact_support": true,
"transport": "smtp",
"proxy_url": "",
"is_minimal": false,
```

```
"issues": [
  {
    "destination": "mailhost",
    "component": "mail_server",
    "node": {
      "name": "node3",
      "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
        }
      }
    },
    "issue": {
      "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'",
      "code": "53149746"
    },
    "corrective_action": {
      "message": "Check the hostname of the SMTP server",
      "code": "53149746"
    }
  },
  {
    "destination": "https://support.netapp.com/aods/asupmessage",
    "component": "ondemand_server",
    "node": {
      "name": "node3",
      "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
        }
      }
    },
    "issue": {
      "message": "AutoSupport OnDemand is disabled when \"-transport\" is
not set to \"https\".",
      "code": "53149740"
    },
    "corrective_action": {
      "message": "Run \"system node autosupport modify -transport https
-node <node name>\" to set \"-transport\" to \"https\".",
      "code": "53149740"
```

```
      }
    }
  ],
  "smtp_encryption": "none",
  "ondemand_enabled": true,
  "_links": {
    "self": {
      "href": "/api/support/autosupport"
    }
  }
}
```


---

### Retrieve the AutoSupport configuration

GET `/support/autosupport`

**Introduced In:** 9.6

Retrieves the AutoSupport configuration of the cluster and if requested, returns connectivity issues with the AutoSupport configuration.

Important note:

- The **issues** field consists of a list of objects containing details of the node that has a connectivity issue, the issue description, and corrective action you can take to address the issue. When not empty, this indicates a connection issue to the **HTTPS**, **SMTP**, or **AutoSupport On Demand** server.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See Requesting specific fields to learn more.

- `issues`

#### Related ONTAP commands

- `system node autosupport show -instance`
- `system node autosupport check show-details`

#### Learn more

- DOC /support/autosupport

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| contact_support | boolean | Specifies whether to send the AutoSupport messages to vendor support. |
| enabled | boolean | Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off. |
| from | string | The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.<br>• example: postmaster@example.com<br>• format: email<br>• Introduced in: 9.6<br>• x-nullable: true |
| is_minimal | boolean | Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics. |
| issues | array[autosupport_issues] | A list of connectivity issues to the HTTPS/SMTP/AOD AutoSupport destinations on the nodes in the cluster along with the corrective actions. |

| Name | Type | Description |
| --- | --- | --- |
| mail_hosts | array[string] | The names of the mail servers used to deliver AutoSupport messages via SMTP. |
| ondemand_enabled | boolean | Specifies whether the AutoSupport OnDemand feature is enabled. When AutoSupport OnDemand is enabled, support personnel can remotely trigger new AutoSupport messages, resend existing AutoSupport messages, and decline the delivery of unwanted AutoSupport messages. When this option is disabled, the cluster does not respond to any AutoSupport OnDemand requests from support personnel. |
| partner_addresses | array[string] | The list of partner addresses. |
| proxy_url | string | Proxy server for AutoSupport message delivery via HTTPS. Optionally specify a username/password for authentication with the proxy server. |
| smtp_encryption | string | The encryption protocol used to deliver AutoSupport messages via SMTP to the configured mail_hosts. |
| to | array[string] | The e-mail addresses to which the AutoSupport messages are sent. |
| transport | string | The name of the transport protocol used to deliver AutoSupport messages.<br>'http' transport is no longer supported by AutoSupport servers. |

**Example response**


```
{
  "contact_support": 1,
  "enabled": 1,
  "from": "<a href="
mailto:postmaster@example.com">postmaster@example.com</a>",
  "is_minimal": 1,
  "issues": [
    {
      "component": "mail_server",
      "corrective_action": {
        "code": "53149746",
        "message": "Check the hostname of the SMTP server"
      },
      "destination": "mailhost1.example.com",
      "issue": {
        "code": "53149746",
        "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "mail_hosts": [
    "mailhost1.example.com",
    "mailhost2.example.com"
  ],
  "ondemand_enabled": 1,
  "partner_addresses": [
    "<a href="mailto:user1@partner.com">user1@partner.com</a>",
    "<a href="mailto:user2@partner.com">user2@partner.com</a>"
  ],
  "proxy_url": "proxy.company.com",
  "smtp_encryption": "none",
  "to": [
    "<a href="mailto:user1@example.com">user1@example.com</a>",
    "<a href="mailto:user2@example.com">user2@example.com</a>"
```

```
  ],
  "transport": "smtp"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

autosupport_connectivity_corrective_action

| Name | Type | Description |
|---|---|---|
| code | string | Corrective action code |
| message | string | Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI.<br>• example: Check the hostname of the SMTP server<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

autosupport_connectivity_issue

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

autosupport_issues

| Name | Type | Description |
| --- | --- | --- |
| component | string | The name of the component where the issue occurred. |
| corrective_action | autosupport_connectivity_corrective_action | |
| destination | string | The HTTPS/SMTP/AOD AutoSupport Destination. |
| issue | autosupport_connectivity_issue | |
| node | node | |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the AutoSupport configuration for an entire cluster

PATCH `/support/autosupport`

**Introduced In:** 9.6

Updates the AutoSupport configuration for the entire cluster.

#### Related ONTAP commands

- `system node autosupport modify`

#### Learn more

- DOC /support/autosupport

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| force | boolean | query | False | Set the force flag to true to modify some of the AutoSupport configurations that are otherwise blocked when the automatic update feature is enabled. Without this flag set to true, an attempt to disable AutoSupport, modify the transport to SMTP, or disable the AutoSupport OnDemand feature fails if the automatic update feature is enabled.<br>• Introduced in: 9.16<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| contact_support | boolean | Specifies whether to send the AutoSupport messages to vendor support. |
| enabled | boolean | Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off. |

| Name | Type | Description |
| --- | --- | --- |
| from | string | The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.<br>• example: postmaster@example.com<br>• format: email<br>• Introduced in: 9.6<br>• x-nullable: true |
| is_minimal | boolean | Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics. |
| issues | array[autosupport_issues] | A list of connectivity issues to the HTTPS/SMTP/AOD AutoSupport destinations on the nodes in the cluster along with the corrective actions. |
| mail_hosts | array[string] | The names of the mail servers used to deliver AutoSupport messages via SMTP. |
| ondemand_enabled | boolean | Specifies whether the AutoSupport OnDemand feature is enabled. When AutoSupport OnDemand is enabled, support personnel can remotely trigger new AutoSupport messages, resend existing AutoSupport messages, and decline the delivery of unwanted AutoSupport messages. When this option is disabled, the cluster does not respond to any AutoSupport OnDemand requests from support personnel. |
| partner_addresses | array[string] | The list of partner addresses. |

| Name | Type | Description |
| --- | --- | --- |
| proxy_url | string | Proxy server for AutoSupport message delivery via HTTPS. Optionally specify a username/password for authentication with the proxy server. |
| smtp_encryption | string | The encryption protocol used to deliver AutoSupport messages via SMTP to the configured mail_hosts. |
| to | array[string] | The e-mail addresses to which the AutoSupport messages are sent. |
| transport | string | The name of the transport protocol used to deliver AutoSupport messages.<br>'http' transport is no longer supported by AutoSupport servers. |

**Example request**


```
{
  "contact_support": 1,
  "enabled": 1,
  "from": "<a href="
mailto:postmaster@example.com">postmaster@example.com</a>",
  "is_minimal": 1,
  "issues": [
    {
      "component": "mail_server",
      "corrective_action": {
        "code": "53149746",
        "message": "Check the hostname of the SMTP server"
      },
      "destination": "mailhost1.example.com",
      "issue": {
        "code": "53149746",
        "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'"
      },
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "mail_hosts": [
    "mailhost1.example.com",
    "mailhost2.example.com"
  ],
  "ondemand_enabled": 1,
  "partner_addresses": [
    "<a href="mailto:user1@partner.com">user1@partner.com</a>",
    "<a href="mailto:user2@partner.com">user2@partner.com</a>"
  ],
  "proxy_url": "proxy.company.com",
  "smtp_encryption": "none",
  "to": [
    "<a href="mailto:user1@example.com">user1@example.com</a>",
    "<a href="mailto:user2@example.com">user2@example.com</a>"
  ],
  "transport": "smtp"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8650862 | The SMTP mail host provided cannot be empty |
| 8650863 | A maximum of 5 SMTP mail hosts can be provided |
| 8650864 | A maximum of 5 email addresses can be provided |
| 8650865 | A maximum of 5 partner email addresses can be provided |
| 8650880 | Automatic update requires AutoSupport to be enabled |
| 8650881 | Automatic update requires AutoSupport to use the HTTPS transport |
| 8650882 | Automatic update requires AutoSupport OnDemand to be enabled |
| 8650886 | The provided parameter requires an effective cluster version of ONTAP 9.15.1 or later |
| 8650887 | Username or password is not allowed in the AutoSupport url and put-url fields |
| 8650889 | The provided parameter requires an effective cluster version of ONTAP 9.16.1 or later |
| 53149727 | The proxy URI provided is invalid |
| 53149728 | The SMTP mail host URI provided is invalid |
| 53149732 | The proxy URI provided is invalid. IPv6 addresses must be enclosed within square brackets |
| 53149737 | The proxy URI provided specifies an unsupported scheme |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

autosupport_connectivity_corrective_action

| Name | Type | Description |
|---|---|---|
| code | string | Corrective action code |
| message | string | Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI.<br>• example: Check the hostname of the SMTP server<br>• readOnly: 1<br>• Introduced in: 9.6<br>• x-nullable: true |

autosupport_connectivity_issue

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

autosupport_issues

| Name | Type | Description |
|---|---|---|
| component | string | The name of the component where the issue occurred. |

| Name | Type | Description |
| --- | --- | --- |
| corrective_action | autosupport_connectivity_corrective_action | |
| destination | string | The HTTPS/SMTP/AOD AutoSupport Destination. |
| issue | autosupport_connectivity_issue | |
| node | node | |

autosupport

| Name | Type | Description |
| --- | --- | --- |
| contact_support | boolean | Specifies whether to send the AutoSupport messages to vendor support. |
| enabled | boolean | Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off. |
| from | string | The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.<br>• example: postmaster@example.com<br>• format: email<br>• Introduced in: 9.6<br>• x-nullable: true |
| is_minimal | boolean | Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics. |
| issues | array[autosupport_issues] | A list of connectivity issues to the HTTPS/SMTP/AOD AutoSupport destinations on the nodes in the cluster along with the corrective actions. |

| Name | Type | Description |
| --- | --- | --- |
| mail_hosts | array[string] | The names of the mail servers used to deliver AutoSupport messages via SMTP. |
| ondemand_enabled | boolean | Specifies whether the AutoSupport OnDemand feature is enabled. When AutoSupport OnDemand is enabled, support personnel can remotely trigger new AutoSupport messages, resend existing AutoSupport messages, and decline the delivery of unwanted AutoSupport messages. When this option is disabled, the cluster does not respond to any AutoSupport OnDemand requests from support personnel. |
| partner_addresses | array[string] | The list of partner addresses. |
| proxy_url | string | Proxy server for AutoSupport message delivery via HTTPS. Optionally specify a username/password for authentication with the proxy server. |
| smtp_encryption | string | The encryption protocol used to deliver AutoSupport messages via SMTP to the configured mail_hosts. |
| to | array[string] | The e-mail addresses to which the AutoSupport messages are sent. |
| transport | string | The name of the transport protocol used to deliver AutoSupport messages.<br>'http' transport is no longer supported by AutoSupport servers. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage AutoSupport messages

### Manage AutoSupport messages

#### Overview

Use this API to invoke and retrieve AutoSupport messages from the nodes in the cluster.

This API supports POST and GET calls. Use a POST call to invoke AutoSupport and a GET call to retrieve AutoSupport messages.

---

#### Examples

### Invoking an AutoSupport on all nodes in the cluster

The following example invokes an AutoSupport on every node in the cluster.
Note that AutoSupport is invoked on all nodes in the cluster if the `node` parameter is omitted. Also, note that the `subject` line is the same when invoking on all nodes.

By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

```
# The API:
POST /support/autosupport/messages

# The call:
curl -X POST "https://<mgmt-
```

```
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"
message\": \"test_msg\", \"type\": \"all\"}"

# The response:
201 CREATED
{
"num_records": 6,
"records": [
  {
    "index": 4,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/4/http"
      }
    }
  },
  {
    "index": 4,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/4/smtp"
      }
    }
```

```
  },
  {
    "index": 4,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/4/noteto"
      }
    }
  },
  {
    "index": 2,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/2/http"
      }
    }
  },
  {
    "index": 2,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "_links": {
        "self": {
```

```
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/2/smtp"
      }
    }
  },
  {
    "index": 2,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/2/noteto"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages/"
  }
}
}
```


---

#### Invoking an AutoSupport on a single node

The following examples invoke an AutoSupport on a single node in the cluster.
Note that AutoSupport is invoked on all nodes in the cluster if the `node` parameter is omitted. You can specify the node-name with either `node` or `node.name` parameter. You can also specify UUID of the node with the `node.uuid` parameter.

---

By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

```
# The API:
POST /support/autosupport/messages

# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"message\": \"test_msg\", \"type\": \"test\", \"node\": { \"name\":
\"node1\" }}"

# The response:
201 CREATED
{
"num_records": 3,
"records": [
  {
    "index": 8,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/8/http"
      }
    }
  },
  {
    "index": 8,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
```

```
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/8/smtp"
      }
    }
  },
  {
    "index": 8,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/8/noteto"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages/"
  }
}
}
```


```
# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"message\": \"test_msg\", \"type\": \"test\", \"node.name\":\"node2\"}"

# The response:
201 CREATED
```

```
{
"num_records": 3,
"records": [
  {
    "index": 4,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/4/http"
      }
    }
  },
  {
    "index": 4,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/4/smtp"
      }
    }
  },
  {
    "index": 4,
    "node": {
      "name": "node2",
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
```

```
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/4/noteto"
      }
    }
  },
],
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages/"
  }
}
}
```

```
# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"message\": \"test_msg\", \"type\": \"test\", \"node.uuid\":\"092e0298-
f250-11e8-9a05-005056bb6666\"}"

# The response:
201 CREATED
{
"num_records": 3,
"records": [
  {
    "index": 5,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
```

```
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/5/http"
      }
    }
  },
  {
    "index": 5,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/5/smtp"
      }
    }
  },
  {
    "index": 5,
    "node": {
      "name": "node1",
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/5/noteto"
      }
    }
  }
```

```
],
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages/"
  }
}
}
```

#### Retrieving AutoSupport messages from all nodes in the cluster

The following example retrieves AutoSupport messages from every node in the cluster.
Note that if the *fields=** parameter is not specified, only node, index, and destination fields are returned.
Filters can be added on the fields to limit the results.

```
# The API:
GET /support/autosupport/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/autosupport/messages?fields=*&return_timeout=15" -H
"accept: application/hal+json"

# The response:
200 OK
{
"records": [
  {
    "node": {
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "index": 1,
    "destination": "smtp",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
```

```
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/smtp"
      }
    }
  },
  {
    "node": {
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "index": 1,
    "destination": "http",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "sent_successful",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/http"
      }
    }
  },
  {
    "node": {
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "index": 1,
    "destination": "noteto",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
```

```
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/noteto"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "index": 1,
    "destination": "smtp",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/1/smtp"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "index": 1,
    "destination": "http",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "sent_successful",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
```

```
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/1/http"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
        }
      }
    },
    "index": 1,
    "destination": "noteto",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/1/noteto"
      }
    }
  }
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages?fields=*&return_timeout=15"
  }
}
}
```


#### Retrieving AutoSupport messages from a specific node and has 'sent_successful' state

The following example retrieves AutoSupport messages from a specific node in the cluster.
Note that if the `fields=*` parameter is not specified, only node, index, and destination fields are returned.
This example uses a filter on the `node.name` and `state` fields. You can add filters to any fields in the response.

```
# The API:
GET /support/autosupport/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/autosupport/messages?node.name=node1&state=sent_successful
&fields=*&return_timeout=15" -H "accept: application/hal+json"

# The response:
200 OK
{
"records": [
  {
    "node": {
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
        }
      }
    },
    "index": 1,
    "destination": "http",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "sent_successful",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/http"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href":
"/api/support/autosupport/messages?node.name=node1&state=sent_successful&f
ields=*&return_timeout=15"
  }
}
}
```

### Retrieve AutoSupport message history

GET `/support/autosupport/messages`

**Introduced In:** 9.6

Retrieves AutoSupport message history from all nodes in the cluster.

There can be a short delay on invoked AutoSupport messages showing in history, dependent on processing of other AutoSupports in the queue.

#### Related ONTAP commands

- `system node autosupport history show`

#### Learn more

- DOC /support/autosupport/messages

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | query | False | Filter by node.uuid |
| node.name | string | query | False | Filter by node.name |
| destination | string | query | False | Filter by destination |
| state | string | query | False | Filter by state |
| generated_on | string | query | False | Filter by generated_on |
| index | integer | query | False | Filter by index |
| subject | string | query | False | Filter by subject |
| error.message | string | query | False | Filter by error.message |
| error.code | integer | query | False | Filter by error.code |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |
| Specify the fields to return. | max_records | integer | query | False |
| Limit the number of records returned. | return_timeout | integer | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[autosupport_message] | List of messages invoked on the node |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "destination": "http",
      "error": {
        "code": "53149746",
        "message": "Could not resolve host: test.com"
      },
      "generated_on": "2019-03-25 17:30:04 -0400",
      "index": 9,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "sent_successful",
      "subject": "WEEKLY_LOG"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

error

Last error during delivery attempt. Empty if "status=sent-successful".

| Name | Type | Description |
|---|---|---|
| code | integer | Error code |
| message | string | Error message |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

autosupport_message

| Name | Type | Description |
|---|---|---|
| _links | self_link | |

| Name | Type | Description |
|---|---|---|
| destination | string | Destination for the AutoSupport |
| error | error | Last error during delivery attempt. Empty if "status=sent-successful". |
| generated_on | string | Date and Time of AutoSupport generation in ISO-8601 format |
| index | integer | Sequence number of the AutoSupport |
| node | node | |
| state | string | State of AutoSupport delivery |
| subject | string | Subject line for the AutoSupport |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create and send an AutoSupport message**

POST `/support/autosupport/messages`

**Introduced In:** 9.6

Creates and sends an AutoSupport message with the provided input parameters.

Important note:

- By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

#### Recommended optional properties

- `message` - Message included in the AutoSupport subject. Use this to identify the generated AutoSupport message.

#### Default property values

If not specified in POST, the following are the default property values:

- `type` - *all*
- `node.name` or `node.uuid` - Not specifying these properties invokes AutoSupport on all nodes in the cluster.

#### Related ONTAP commands

- `system node autosupport invoke`

#### Learn more

- DOC /support/autosupport/messages

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| destination | string | Destination for the AutoSupport |
| error | error | Last error during delivery attempt. Empty if "status=sent-successful". |
| generated_on | string | Date and Time of AutoSupport generation in ISO-8601 format |
| index | integer | Sequence number of the AutoSupport |

| Name | Type | Description |
|---|---|---|
| message | string | Message included in the AutoSupport subject |
| node | node | |
| state | string | State of AutoSupport delivery |
| subject | string | Subject line for the AutoSupport |
| type | string | Type of AutoSupport collection to issue |
| uri | string | Alternate destination for the AutoSupport |

**Example request**

```
{
  "destination": "http",
  "error": {
    "code": "53149746",
    "message": "Could not resolve host: test.com"
  },
  "generated_on": "2019-03-25 17:30:04 -0400",
  "index": 9,
  "message": "invoked_test_autosupport_rest",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "sent_successful",
  "subject": "WEEKLY_LOG",
  "type": "test",
  "uri": "https://1.2.3.4/delivery_uri"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[autosupport_message] | List of messages invoked on the node |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "destination": "http",
      "error": {
        "code": "53149746",
        "message": "Could not resolve host: test.com"
      },
      "generated_on": "2019-03-25 17:30:04 -0400",
      "index": 9,
      "message": "invoked_test_autosupport_rest",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "sent_successful",
      "subject": "WEEKLY_LOG",
      "type": "test",
      "uri": "https://1.2.3.4/delivery_uri"
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 8650851 | AutoSupport invocation failed because AutoSupport collection is disabled |
| 8650852 | AutoSupport invocation failed because the trigger is disabled |
| 8650853 | AutoSupport invocation failed because the trigger rate is higher than the maximum supported value |
| 8650866 | The message parameter is not supported with performance AutoSupports |
| 8650884 | The specified field is not allowed for the current operation |
| 8650885 | The specified field cannot be empty |
| 53149748 | The destination URI provided for the invoked AutoSupport is invalid |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

error

Last error during delivery attempt. Empty if "status=sent-successful".

| Name | Type | Description |
|---|---|---|
| code | integer | Error code |
| message | string | Error message |

_links

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

autosupport_message

| Name | Type | Description |
|---|---|---|
| destination | string | Destination for the AutoSupport |
| error | error | Last error during delivery attempt. Empty if "status=sent-successful". |
| generated_on | string | Date and Time of AutoSupport generation in ISO-8601 format |
| index | integer | Sequence number of the AutoSupport |
| message | string | Message included in the AutoSupport subject |
| node | node | |
| state | string | State of AutoSupport delivery |

| Name | Type | Description |
|---|---|---|
| subject | string | Subject line for the AutoSupport |
| type | string | Type of AutoSupport collection to issue |
| uri | string | Alternate destination for the AutoSupport |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Retrieve information about an AutoSupport message**

GET `/support/autosupport/messages/{node.uuid}/{index}/{destination}`

**Introduced In:** 9.14

Retrieves information about a single Autosupport message.

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | string | path | True | AutoSupport message sequence number |
| destination | string | path | True | The destination for this AutoSupport message.<br><br>• enum: ["smtp", "http", "noteto", "retransmit"] |
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| destination | string | Destination for the AutoSupport |
| error | error | Last error during delivery attempt. Empty if "status=sent-successful". |
| generated_on | string | Date and Time of AutoSupport generation in ISO-8601 format |
| index | integer | Sequence number of the AutoSupport |
| node | node | |
| state | string | State of AutoSupport delivery |
| subject | string | Subject line for the AutoSupport |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "destination": "http",
  "error": {
    "code": "53149746",
    "message": "Could not resolve host: test.com"
  },
  "generated_on": "2019-03-25 17:30:04 -0400",
  "index": 9,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "sent_successful",
  "subject": "WEEKLY_LOG"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

error

Last error during delivery attempt. Empty if "status=sent-successful".

| Name | Type | Description |
|---|---|---|
| code | integer | Error code |
| message | string | Error message |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster configuration backup information

### Manage cluster configuration backup information

#### Overview

This API retrieves the current settings for the configuration and updates configuration backup settings. The GET operation retrieves the current settings for the configuration and the PATCH operation updates the configuration backup settings.

#### Examples

These examples show how to retrieve and update the configuration backup settings.

### Retrieving the configuration backup settings

```
# The API:
/api/support/configuration-backup

# The call:
curl -X GET "https://<mgmt-ip>/api/support/configuration-backup" -H
"accept: application/hal+json"

# The response:
{
  "url": "http://10.224.65.198/backups",
  "username": "me",
  "_links": {
      "self": {
          "href": "/api/support/configuration-backup"
      }
  }
}
```

### Updating the configuration backup settings

```
# The API:
/api/support/configuration-backup

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/configuration-backup" -H
"accept: application/hal+json"

# The body:
{
  "url": "https://10.224.65.198/new_backups",
  "username": "new_me",
  "password": "new_pass",
  "validate_certificate": "true",
  "rest_method": "post"
}

# The response header:
HTTP/1.1 200 OK
Date: Tue, 05 Jun 2018 18:17:48 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/hal+json
```

### Retrieve cluster configuration backup information

GET `/support/configuration-backup`

**Introduced In:** 9.6

Retrieves the cluster configuration backup information.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

**Example response**

```
{
  "password": "yourpassword",
  "rest_method": "post",
  "url": "http://10.224.65.198/backups",
  "username": "me"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update cluster configuration backup information

PATCH `/support/configuration-backup`

**Introduced In:** 9.6

Updates the cluster configuration backup information.

#### Request Body

| Name | Type | Description |
|---|---|---|
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

#### Example request

```
{
  "password": "yourpassword",
  "rest_method": "post",
  "url": "http://10.224.65.198/backups",
  "username": "me"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 39387137 | Invalid URL |
| 39387138 | Invalid URL Scheme |
| 5963792 | Using HTTPS for URL requires an effective cluster version of 9.7.0 or later. |
| 5963794 | Validating the certificate is supported for the HTTPS protocol only. |

**Definitions**

**See Definitions**

configuration_backup

| Name | Type | Description |
|---|---|---|
| password | string | |
| rest_method | string | The REST API HTTP method (POST/PUT). |
| url | string | An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster. |
| username | string | |
| validate_certificate | boolean | Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage configuration backups

### Retrieve configuration backup files

GET `/support/configuration-backup/backups`

**Introduced In:** 9.7

Retrieves a list of configuration backup files.

#### Related ONTAP commands

- `system configuration backup show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| type | string | query | False | Filter by type |
| name | string | query | False | Filter by name |
| size | integer | query | False | Filter by size |
| auto | boolean | query | False | Filter by auto |
| time | string | query | False | Filter by time |
| version | string | query | False | Filter by version |
| backup_nodes.name | string | query | False | Filter by backup_nodes.name |
| download_link | string | query | False | Filter by download_link |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |
| desc] direction. Default direction is 'asc' for ascending. | fields | array[string] | query | False |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[configuration_backup_file] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "backup_nodes": [
        {
          "name": "string"
        }
      ],
      "download_link": "https://10.224.65.198/backups/backup_file.7z",
      "name": "backup_file.7z",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 4787563,
      "time": "2019-02-04 13:33:48 -0500",
      "type": "string",
      "version": "9.7.0"
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

backup_node

The node name in the backup.

| Name | Type | Description |
|---|---|---|
| name | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

The node that owns the configuration backup.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

configuration_backup_file

The configuration backup file.

| Name | Type | Description |
|---|---|---|
| auto | boolean | Indicates if the backup was created automatically. |
| backup_nodes | array[backup_node] | The list of nodes included in the backup. |

| Name | Type | Description |
|---|---|---|
| download_link | string | The link to download the backup file. |
| name | string | The backup name. |
| node | node | The node that owns the configuration backup. |
| size | integer | The size of the backup in bytes. |
| time | string | The backup creation time. |
| type | string | The backup type. |
| version | string | The software version. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create a configuration backup**

POST `/support/configuration-backup/backups`

**Introduced In:** 9.7

Creates a configuration backup. The required backup file name must end with .7z extension.

#### Required properties

- `node.uuid` or `node.name` - The node UUID or node name on which the configuration backup will be created.
- `name` - The backup file name

#### Related ONTAP commands

- `system configuration backup create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| auto | boolean | Indicates if the backup was created automatically. |
| backup_nodes | array[backup_node] | The list of nodes included in the backup. |
| download_link | string | The link to download the backup file. |
| name | string | The backup name. |
| node | node | The node that owns the configuration backup. |
| size | integer | The size of the backup in bytes. |
| time | string | The backup creation time. |
| type | string | The backup type. |
| version | string | The software version. |

### Example request

```
{
  "backup_nodes": [
    {
      "name": "string"
    }
  ],
  "download_link": "https://10.224.65.198/backups/backup_file.7z",
  "name": "backup_file.7z",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 4787563,
  "time": "2019-02-04 13:33:48 -0500",
  "type": "string",
  "version": "9.7.0"
}
```

#### Response

```
Status: 202, Accepted
```

| Name | Type | Description |
|---|---|---|
| job | job_link | |

### Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Response

```
Status: 201, Created
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5963818 | Failed to queue private job. |
| 5963819 | Failed to queue cluster job. |
| 5963871 | Requested backup name is missing .7z suffix. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

backup_node

The node name in the backup.

| Name | Type | Description |
|---|---|---|
| name | string | |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

node

The node that owns the configuration backup.

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

configuration_backup_file

The configuration backup file.

| Name | Type | Description |
|---|---|---|
| auto | boolean | Indicates if the backup was created automatically. |
| backup_nodes | array[backup_node] | The list of nodes included in the backup. |
| download_link | string | The link to download the backup file. |
| name | string | The backup name. |
| node | node | The node that owns the configuration backup. |
| size | integer | The size of the backup in bytes. |
| time | string | The backup creation time. |

| Name | Type | Description |
| --- | --- | --- |
| type | string | The backup type. |
| version | string | The software version. |

job_link

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a configuration backup

DELETE `/support/configuration-backup/backups/{node.uuid}/{name}`

**Introduced In:** 9.7

Deletes a configuration backup.

#### Related ONTAP commands

- `system configuration backup delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | UUID of the node that owns the configuration backup. |
| name | string | path | True | Name of the configuration backup to be deleted. |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5963826 | Failed to delete backup file. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve configuration backup file details

GET `/support/configuration-backup/backups/{node.uuid}/{name}`

**Introduced In:** 9.7

Retrieves details of the specified configuration backup file.

#### Related ONTAP commands

- `system configuration backup show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | UUID of the node that owns the configuration backup. |
| name | string | path | True | Name of the configuration backup. |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| auto | boolean | Indicates if the backup was created automatically. |
| backup_nodes | array[backup_node] | The list of nodes included in the backup. |
| download_link | string | The link to download the backup file. |
| name | string | The backup name. |
| node | node | The node that owns the configuration backup. |
| size | integer | The size of the backup in bytes. |
| time | string | The backup creation time. |

| Name | Type | Description |
|---|---|---|
| type | string | The backup type. |
| version | string | The software version. |

**Example response**

```
{
  "backup_nodes": [
    {
      "name": "string"
    }
  ],
  "download_link": "https://10.224.65.198/backups/backup_file.7z",
  "name": "backup_file.7z",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 4787563,
  "time": "2019-02-04 13:33:48 -0500",
  "type": "string",
  "version": "9.7.0"
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 5963777 | Configuration backup file does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

backup_node

The node name in the backup.

| Name | Type | Description |
|---|---|---|
| name | string | |

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

The node that owns the configuration backup.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage core dumps

### Manage core dumps

#### Overview

The core dump GET API retrieves all of the core dumps on the cluster or a node. The core dump DELETE API deletes a specified core dump.
Some fields are only populated for core dump type "kernel". Refer to the model for further information. Fields will not be displayed if they are not populated.
A core can be deleted even if the core is in the process of being saved.

#### Examples

### 1) Retrieving a list of core dumps from the cluster

The following example returns a list of core dumps on the cluster:

```
# The API:
/api/support/coredump/coredumps

# The call:
curl -X GET "https://<mgmt-ip>/api/support/coredump/coredumps" -H "accept:
application/hal+json"

# The response:
{
"records": [
  {
    "node": {
      "uuid": "227683c1-e9c7-11eb-b995-005056bbbfb3",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/227683c1-e9c7-11eb-b995-
005056bbbfb3"
        }
      }
    },
    "type": "kernel",
    "name": "core.4136886422.2021-07-21.20_20_53.nz",
```

```
    "_links": {
      "self": {
        "href": "/api/support/coredump/coredumps/227683c1-e9c7-11eb-b995-
005056bbbfb3/core.4136886422.2021-07-21.20_20_53.nz"
      }
    }
  },
  {
    "node": {
      "uuid": "227683c1-e9c7-11eb-b995-005056bbbfb3",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/227683c1-e9c7-11eb-b995-
005056bbbfb3"
        }
      }
    },
    "type": "application",
    "name": "mlogd.968.4136886422.2021-07-22.01_10_01.ucore.bz2",
    "_links": {
      "self": {
        "href": "/api/support/coredump/coredumps/227683c1-e9c7-11eb-b995-
005056bbbfb3/mlogd.968.4136886422.2021-07-22.01_10_01.ucore.bz2"
      }
    }
  },
  {
    "node": {
      "uuid": "d583d44e-e9c6-11eb-a270-005056bb47f9",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/d583d44e-e9c6-11eb-a270-
005056bb47f9"
        }
      }
    },
    "type": "kernel",
    "name": "core.4136886421.2021-07-21.17_57_02.nz",
    "_links": {
      "self": {
        "href": "/api/support/coredump/coredumps/d583d44e-e9c6-11eb-a270-
005056bb47f9/core.4136886421.2021-07-21.17_57_02.nz"
      }
    }
```

```
  },
  {
    "node": {
      "uuid": "d583d44e-e9c6-11eb-a270-005056bb47f9",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/d583d44e-e9c6-11eb-a270-
005056bb47f9"
        }
      }
    },
    "type": "application",
    "name": "mlogd.979.4136886421.2021-07-22.01_11_37.ucore.bz2",
    "_links": {
      "self": {
        "href": "/api/support/coredump/coredumps/d583d44e-e9c6-11eb-a270-
005056bb47f9/mlogd.979.4136886421.2021-07-22.01_11_37.ucore.bz2"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/support/coredump/coredumps"
  }
}
}
```


**2) Retrieving a specific core dump**

The following example returns the requested core dump. If there is no core dump with the requested node UUID and name, an error is returned.

```
# The API:
/api/support/coredump/coredumps/{node.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/coredump/coredumps/d583d44e-
e9c6-11eb-a270-005056bb47f9/core.4136886421.2021-07-21.17_57_02.nz" -H
"accept: application/hal+json"

# The response:
{
"node": {
  "uuid": "d583d44e-e9c6-11eb-a270-005056bb47f9",
  "name": "node1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/d583d44e-e9c6-11eb-a270-005056bb47f9"
    }
  }
},
"type": "kernel",
"name": "core.4136886421.2021-07-21.17_57_02.nz",
"panic_time": "2021-07-21T13:57:02-04:00",
"size": 945111148,
"md5_data_chksum": "611a6fa50d230a1af402f3724ab320af",
"is_saved": true,
"_links": {
  "self": {
    "href": "/api/support/coredump/coredumps/d583d44e-e9c6-11eb-a270-
005056bb47f9/core.4136886421.2021-07-21.17_57_02.nz"
  }
}
}
```

#### 3) Deleting a core dump

The following example deletes the requested core dump. If there is no core dump with the requested node UUID and name to delete, an error is returned.

```
# The API:
/api/support/coredump/coredumps/{node.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/coredump/coredumps/d583d44e-
e9c6-11eb-a270-005056bb47f9/mlogd.979.4136886421.2021-07-
22.01_11_37.ucore.bz2" -H "accept: application/hal+json"

# The response header:
HTTP/1.1 200 OK
Date: Thu, 22 Jul 2021 01:32:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/hal+json
```

---

### Retrieve core dumps

GET `/support/coredump/coredumps`

**Introduced In:** 9.10

Retrieves a collection of coredumps.

#### Related ONTAP commands

- `system node coredump show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| type | string | query | False | Filter by type |
| is_partial | boolean | query | False | Filter by is_partial |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| md5_data_checksum | string | query | False | Filter by md5_data_checksum |
| is_saved | boolean | query | False | Filter by is_saved |
| panic_time | string | query | False | Filter by panic_time |
| size | integer | query | False | Filter by size |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[coredump] | |

#### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "md5_data_checksum": "5118488cc5065e33a16986001b1ffa48",
      "name": "core.4136886413.2021-03-01.22_09_11.nz",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "panic_time": "2021-03-01 18:09:11 -0400",
      "size": 1161629804,
      "type": "kernel"
    }
  ]
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

coredump

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| is_partial | boolean | Specifies whether or not the core is a partial core. Applicable only to kernel core dump type. |
| is_saved | boolean | Specifies whether or not the core file is saved. |
| md5_data_checksum | string | MD5 checksum of the compressed data of core. Applicable only to kernel core dump type. |
| name | string | Core name |
| node | node | |

| Name | Type | Description |
|---|---|---|
| panic_time | string | Time of panic that generated the core. |
| size | integer | Size of core, in bytes. Applicable only to kernel core dump type. |
| type | string | Core type, kernel or application |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Delete a core dump

DELETE `/support/coredump/coredumps/{node.uuid}/{name}`

**Introduced In:** 9.10

Deletes a core dump.

#### Related ONTAP commands

- `system node coredump delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| name | string | path | True | Core name |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve a core dump

GET `/support/coredump/coredumps/{node.uuid}/{name}`

**Introduced In:** 9.10

Retrieves a specific core dump.

#### Related ONTAP commands

- `system node coredump show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| node.uuid | string | path | True | Node UUID |
| name | string | path | True | Core name |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| is_partial | boolean | Specifies whether or not the core is a partial core. Applicable only to kernel core dump type. |
| is_saved | boolean | Specifies whether or not the core file is saved. |
| md5_data_checksum | string | MD5 checksum of the compressed data of core. Applicable only to kernel core dump type. |
| name | string | Core name |
| node | node | |
| panic_time | string | Time of panic that generated the core. |

| Name | Type | Description |
|---|---|---|
| size | integer | Size of core, in bytes. Applicable only to kernel core dump type. |
| type | string | Core type, kernel or application |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "md5_data_checksum": "5118488cc5065e33a16986001b1ffa48",
  "name": "core.4136886413.2021-03-01.22_09_11.nz",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "panic_time": "2021-03-01 18:09:11 -0400",
  "size": 1161629804,
  "type": "kernel"
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage EMS configuration

### Manage EMS configuration

#### Overview

This endpoint is used to configure general parameters of the Event Management System (EMS).

#### Examples

### Configuring the system-wide email parameters

```
# The API:
PATCH /support/ems

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/ems" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_config_patch.txt"
test_ems_config_patch.txt(body):
{
"mail_from": "administrator@mycompany.com",
"mail_server": "mycompany.com",
"mail_server_user": "smtp"
}

# The response:
200 OK
```

### Retrieving the EMS configuration

The following example retrieves EMS configuration for the cluster.

```
# The API:
GET /support/ems

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems" -H "accept:
application/hal+json"

# The response:
200 OK
{
"mail_from": "admin@localhost",
"mail_server": "localhost",
"pubsub_enabled": true,
"_links": {
  "self": {
    "href": "/api/support/ems"
  }
}
}
```

### Retrieve an EMS configuration

GET `/support/ems`

**Introduced In:** 9.6

Retrieves the EMS configuration.

#### Related ONTAP commands

- `event config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| mail_from | string | Mail from |
| mail_server | string | Mail server (SMTP) |
| mail_server_password | string | Password for Mail server (SMTP) |
| mail_server_user | string | Username for Mail server (SMTP) |
| proxy_password | string | Password for HTTP/HTTPS proxy |
| proxy_url | string | HTTP/HTTPS proxy URL |
| proxy_user | string | User name for HTTP/HTTPS proxy |
| pubsub_enabled | boolean | Is Publish/Subscribe Messaging Enabled? |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "mail_from": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "mail_server": "mail.mycompany.com",
  "mail_server_password": "password",
  "mail_server_user": "user",
  "proxy_password": "password",
  "proxy_url": "https://proxyserver.mycompany.com",
  "proxy_user": "proxy_user",
  "pubsub_enabled": 1
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an EMS configuration

PATCH `/support/ems`

**Introduced In:** 9.6

Updates the EMS configuration.

#### Related ONTAP commands

- `event config modify`

#### Request Body

| Name | Type | Description |
|---|---|---|
| mail_from | string | Mail from |
| mail_server | string | Mail server (SMTP) |
| mail_server_password | string | Password for Mail server (SMTP) |
| mail_server_user | string | Username for Mail server (SMTP) |
| proxy_password | string | Password for HTTP/HTTPS proxy |
| proxy_url | string | HTTP/HTTPS proxy URL |
| proxy_user | string | User name for HTTP/HTTPS proxy |
| pubsub_enabled | boolean | Is Publish/Subscribe Messaging Enabled? |

#### Example request

```
{
  "mail_from": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "mail_server": "mail.mycompany.com",
  "mail_server_password": "password",
  "mail_server_user": "user",
  "proxy_password": "password",
  "proxy_url": "https://proxyserver.mycompany.com",
  "proxy_user": "proxy_user",
  "pubsub_enabled": 1
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983123 | The validation of the mail server provided failed |
| 983137 | The proxy URL provided is invalid |
| 983139 | The IPv6 proxy URL provided is invalid |
| 983140 | The proxy URL provided contains an invalid scheme. The only supported scheme is 'http' |
| 983160 | The provided parameter requires an effective cluster version of ONTAP 9.9.1 or later |
| 983220 | The proxy URL provided contains an invalid scheme. The only supported schemes are 'http' and 'https' |
| 983224 | The mail server URL cannot contain a username or password |
| 983225 | The mail server username cannot be set without a mail server URL |
| 983226 | The mail server password cannot be set without a mail server username |
| 983227 | The provided field requires an effective cluster version of ONTAP 9.15.1 or later |
| 983229 | The proxy URL cannot contain a username or password |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

ems_config

| Name | Type | Description |
|---|---|---|
| mail_from | string | Mail from |
| mail_server | string | Mail server (SMTP) |
| mail_server_password | string | Password for Mail server (SMTP) |
| mail_server_user | string | Username for Mail server (SMTP) |
| proxy_password | string | Password for HTTP/HTTPS proxy |
| proxy_url | string | HTTP/HTTPS proxy URL |
| proxy_user | string | User name for HTTP/HTTPS proxy |
| pubsub_enabled | boolean | Is Publish/Subscribe Messaging Enabled? |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage application logs

### Manage application logs

#### Overview

A POST request allows a standalone application to log its events on the ONTAP system as EMS events and optionally generate AutoSupports. This event information is encapsulated in an app.log.x event, or callhome.client.app.x event in case AutoSupport generation is required, based on the event's severity.

> This API only supports POST. Generated events can be retrieved using the /api/support/ems/events API.

#### Examples

### Creating a new app.log.* event

The following example generates an app.log.error event.

```
# The API:
POST /api/support/ems/application-logs

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/application-logs" -H
"accept: application/hal+json" -H "Content-Type: application/json" -d "{
\"computer_name\": \"W12-R2-DC\", \"event_source\": \"SnapCenter
Software\", \"app_version\": \"3.0.0.2373\", \"event_id\": 10001,
\"category\": \"error\", \"event_description\": \"Clone split job
failed\", \"severity\": \"error\", \"autosupport_required\": false}"

# The response:
201 CREATED
{}
```

### Retrieving a list of events generated by this API

The following example retrieves all app.log.* events.

```
# The API:
GET /api/support/ems/events

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/ems/events?message.name=app.log.*" -H "accept:
application/hal+json"

# The response:
200 OK
{
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "d96d5d8d-118a-11ec-aad8-005056bb974c",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/d96d5d8d-118a-11ec-aad8-
005056bb974c"
        }
      }
    },
    "index": 7865,
    "time": "2021-09-17T09:36:28-04:00",
    "message": {
      "severity": "emergency",
      "name": "app.log.emerg"
    },
    "log_message": "app.log.emerg: iMac: iPad ipad3: (10001) handheld:
TestLine1",
    "_links": {
      "self": {
        "href": "/api/support/ems/events/node1/7865"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/support/ems/events?message.name=app.log.*"
  }
}
}
```

### Create an application log

POST `/support/ems/application-logs`

**Introduced In:** 9.11

Creates an app.log.* event. Setting the "autosupport_required" flag causes callhome.client.app.* events also to be generated, which in turn triggers AutoSupports.

#### Required properties

- `computer_name` - Client computer connected to the cluster.
- `event_source` - Client application that generated this event.
- `app_version` - Client application version.
- `event_id` - Application eventID.
- `category` - Event category.
- `event_description` - Event description.
- `severity` - Severity of the event.
- `autosupport_required` - Indicates whether AutoSupport generation is required.

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| app_version | string | Client application version. |
| autosupport_required | boolean | Indicates whether AutoSupport generation is required. |
| category | string | Event category. |
| computer_name | string | Client computer connected to the cluster. |

| Name | Type | Description |
| --- | --- | --- |
| event_description | string | Event description, which corresponds to the subject parameter in the generated event. |
| event_id | integer | Application eventID. |
| event_source | string | Client application that generated this event. |
| severity | string | Severity |

**Example request**

```
{
  "app_version": "3.0.0.2373",
  "autosupport_required": "",
  "category": "Error",
  "computer_name": "W12-R2-DC",
  "event_description": "Clone split job failed",
  "event_id": 10001,
  "event_source": "SnapCenter Software",
  "severity": "error"
}
```

**Response**

```
Status: 201, Created
```

**Headers**

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 983173 | Application log event generation failed. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

ems_application_log

Input for generation of app.log.x and callhome.client.app.x events.

| Name | Type | Description |
|---|---|---|
| app_version | string | Client application version. |
| autosupport_required | boolean | Indicates whether AutoSupport generation is required. |
| category | string | Event category. |
| computer_name | string | Client computer connected to the cluster. |
| event_description | string | Event description, which corresponds to the subject parameter in the generated event. |
| event_id | integer | Application eventID. |
| event_source | string | Client application that generated this event. |
| severity | string | Severity |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

## Manage event destinations

### Manage event destinations

#### Overview

Manages the list of destinations. A destination is defined by a type and a place to which an event's information is transmitted.

#### Destination Types

An EMS destination is defined by a single type, which is one of the following:

- email
- syslog
- rest_api
- snmp

### Email

The 'email' type allows you to define a mailbox where information about an observed event is sent by SMTP. Enter the address in the destination field in a valid format.
For example: administrator@mycompany.com

### Syslog

The 'syslog' type allows you to specify a remote syslog server that can receive information about an observed event. Enter the hostname or IP address in the destination field.
For example: syslog.mycompany.com, 192.168.1.1

You can optionally specify the port and transport protocol to be used. The supported transport protocols are:

- udp_unencrypted
- tcp_unencrypted
- tcp_encrypted

#

If transport is specified and the port is not specified, the default port for each transport is automatically selected. The default port for udp_unencrypted is 514, for tcp_unencrypted 601, and for tcp_encrypted 6514. The default transport protocol is udp_unencrypted.

#

The message format to be used can also be optionally specified. The supported message formats are:

- legacy_netapp (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG)
- rfc_5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG)

#

If the default timestamp or hostname formats in the message format needs to be overridden, use timestamp_override and hostname_override properties.

The supported timestamp formats are:

- no_override (timestamp format based on the syslog.format.message property i.e. rfc_3164 if syslog.format.message is legacy_netapp, iso_8601_local_time if syslog.format.message is rfc_5424)
- rfc_3164 (format: Mmm dd hh:mm:ss)
- iso_8601_local_time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm)
- iso_8601_utc (format: YYYY-MM-DDThh:mm:ssZ)

#

The supported hostname formats are:

- no_override (hostname format based on the syslog.format.message property i.e. fqdn if syslog.format.message is rfc_5424, hostname_only if syslog.format.message is legacy_netapp)
- fqdn (Fully Qualified Domain Name)
- hostname_only

**Rest API**

The 'rest_api' type allows you to define a URL where information about an observed event is sent using the REST protocol. Enter the URL in the destination field.
The URL must contain a valid transmission schema which can be one of the following:

- http
- https

#

Using the *https* schema, you can configure a client-side certificate for mutual authentication. A client-side certificate is specified by the `ca` and `serial_number` fields in the `certificate` object. The `name` field of the `certificate` object is read-only and cannot be used to configure a client-side certificate.

For example: http://rest.mycompany.com, https://192.168.1.1

**SNMP**

The 'snmp' type describes addresses where information about an observed event is sent using SNMP traps. The system defines a default instance of this type and it is restricted to read-only.
This type has the following limitations:

- Cannot create new destinations of type 'snmp'
- Cannot modify the default 'snmp' destination

SNMP trap host details need to be configured through one of the following:

| Type | Command / API |
|---|---|
| CLI | 'system snmp traphost' |
| ZAPI | 'snmp-traphost-add' / 'snmp-traphost-delete' |

#### Examples

### Retrieving the list of active destinations

```
# The API:
GET /api/support/ems/destinations

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/destinations" -H "accept:
application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "name": "snmp-traphost",
    "_links": {
      "self": {
        "href": "/api/support/ems/destinations/snmp-traphost"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/support/ems/destinations"
  }
}
}
```

### Creating a new 'email' destination

```
# The API:
POST /support/ems/destinations

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/destinations" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_destinations_post.txt"
test_ems_destinations_post.txt(body):

# JSON Body
{
"name": "Technician_Email",
"type": "email",
"destination": "technician@mycompany.com",
"filters": [
  { "name" : "critical-wafl" }
]
}

# The response:
201 Created
```

**Creating a new 'syslog' destination**

```
# The API:
POST /support/ems/destinations

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/destinations" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_destinations_syslog_post.txt"
test_ems_destinations_syslog_post.txt(body):

# JSON Body
{
"name": "Syslog_dest",
"type": "syslog",
"destination": "syslog.mycompany.com",
"syslog": {
  "port": 601,
  "transport": "tcp_unencrypted",
  "format": {
    "message": "rfc_5424",
    "hostname_override": "fqdn",
    "timestamp_override": "iso_8601_local_time"
  }
},
"filters": [
  { "name" : "critical-wafl" }
]
}

# The response:
201 Created
```

**Retrieve a collection of event destinations**

GET `/support/ems/destinations`

**Introduced In:** 9.6

Retrieves a collection of event destinations.

#### Related ONTAP commands

- `event notification destination show`
- `event notification show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| system_defined | boolean | query | False | Filter by system_defined<br>• Introduced in: 9.10 |
| destination | string | query | False | Filter by destination |
| certificate.name | string | query | False | Filter by certificate.name<br>• Introduced in: 9.11 |
| certificate.serial_number | string | query | False | Filter by certificate.serial_number<br>• maxLength: 40<br>• minLength: 1 |
| certificate.ca | string | query | False | Filter by certificate.ca<br>• maxLength: 256<br>• minLength: 1 |
| connectivity.errors.message.message | string | query | False | Filter by connectivity.errors.message.message<br>• Introduced in: 9.11 |
| connectivity.errors.message.code | string | query | False | Filter by connectivity.errors.message.code<br>• Introduced in: 9.11 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| connectivity.errors.message.arguments.message | string | query | False | Filter by connectivity.errors.message.arguments.message<br><br>• Introduced in: 9.11 |
| connectivity.errors.message.arguments.code | string | query | False | Filter by connectivity.errors.message.arguments.code<br><br>• Introduced in: 9.11 |
| connectivity.errors.node.name | string | query | False | Filter by connectivity.errors.node.name<br><br>• Introduced in: 9.11 |
| connectivity.errors.node.uuid | string | query | False | Filter by connectivity.errors.node.uuid<br><br>• Introduced in: 9.11 |
| connectivity.state | string | query | False | Filter by connectivity.state<br><br>• Introduced in: 9.11 |
| type | string | query | False | Filter by type |
| syslog.port | integer | query | False | Filter by syslog.port<br><br>• Introduced in: 9.12 |
| syslog.format.hostname_override | string | query | False | Filter by syslog.format.hostname_override<br><br>• Introduced in: 9.12 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| syslog.format.message | string | query | False | Filter by syslog.format.message<br>• Introduced in: 9.12 |
| syslog.format.timestamp_override | string | query | False | Filter by syslog.format.timestamp_override<br>• Introduced in: 9.12 |
| syslog.transport | string | query | False | Filter by syslog.transport<br>• Introduced in: 9.12 |
| access_control_role.name | string | query | False | Filter by access_control_role.name<br>• Introduced in: 9.13 |
| filters.name | string | query | False | Filter by filters.name |
| name | string | query | False | Filter by name |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_control_role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "certificate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ca": "VeriSign",
        "name": "string",
        "serial_number": 1234567890
      },
      "connectivity": {
        "errors": [
          {
            "message": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
```

```
              ],
              "code": "4",
              "message": "entry doesn't exist"
            },
            "node": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ],
        "state": "fail"
      },
      "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
      "filters": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "important-events"
        }
      ],
      "name": "Admin_Email",
      "syslog": {
        "format": {
          "hostname_override": "string",
          "message": "string",
          "timestamp_override": "string"
        },
        "port": 514,
        "transport": "string"
      },
      "system_defined": 1,
      "type": "email"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

certificate

Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ca | string | Client certificate issuing CA |
| name | string | Certificate name |
| serial_number | string | Client certificate serial number |

arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

message

Information to be displayed to the user.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[arguments] | Message arguments |
| code | string | Unique message code. |
| message | string | User message. |

node

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |
| uuid | string | |

errors

Error object included in the event of connectivity failure.

| Name | Type | Description |
| --- | --- | --- |
| message | message | Information to be displayed to the user. |
| node | node | |

connectivity

| Name | Type | Description |
| --- | --- | --- |
| errors | array[errors] | A list of errors encountered during connectivity checks. |
| state | string | Current connectivity state. |

filters

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| name | string | |

format

| Name | Type | Description |
| --- | --- | --- |
| hostname_override | string | Syslog Hostname Format Override. The supported hostname formats are no_override (hostname format based on the syslog.format.message property i.e. fqdn if syslog.format.message is rfc_5424, hostname_only if syslog.format.message is legacy_netapp), fqdn (Fully Qualified Domain Name) and hostname_only. |
| message | string | Syslog Message Format. The supported message formats are legacy_netapp (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG) and rfc_5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG). |
| timestamp_override | string | Syslog Timestamp Format Override. The supported timestamp formats are no_override (timestamp format based on the syslog.format.message property i.e. rfc_3164 if syslog.format.message is legacy_netapp, iso_8601_local_time if syslog.format.message is rfc_5424), rfc_3164 (format: Mmm dd hh:mm:ss), iso_8601_local_time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm) and iso_8601_utc (format: YYYY-MM-DDThh:mm:ssZ). |

syslog

| Name | Type | Description |
|---|---|---|
| format | format | |
| port | integer | Syslog Port. |
| transport | string | Syslog Transport Protocol. |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| name | string | Destination name. Valid in POST. |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |

| Name | Type | Description |
|---|---|---|
| type | string | Type of destination. Valid in POST. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an event destination

POST `/support/ems/destinations`

**Introduced In:** 9.6

Creates an event destination.

#### Required properties

- `name` - String that uniquely identifies the destination.
- `type` - Type of destination that is to be created.
- `destination` - String that identifies the destination. The contents of this property changes depending on type.

#### Recommended optional properties

- `filters.name` - List of filter names that should direct to this destination.
- `certificate` - When specifying a rest api destination, a client certificate can be provided.
- `syslog` - When specifying a syslog destination, a port, transport protocol, message format, timestamp format and hostname format can be provided.

#### Related ONTAP commands

- `event notification destination create`
- `event notification create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| name | string | Destination name. Valid in POST. |
| syslog | syslog | |

| Name | Type | Description |
| --- | --- | --- |
| system_defined | boolean | Flag indicating system-defined destinations. |
| type | string | Type of destination. Valid in POST. |

**Example request**

```
{
  "access_control_role": {
    "name": "admin"
  },
  "certificate": {
    "ca": "VeriSign",
    "name": "string",
    "serial_number": 1234567890
  },
  "connectivity": {
    "errors": [
      {
        "message": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "state": "fail"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ],
  "name": "Admin_Email",
  "syslog": {
    "format": {
      "hostname_override": "string",
      "message": "string",
      "timestamp_override": "string"
    },
```

```
    "port": 514,
    "transport": "string"
  },
  "system_defined": 1,
  "type": "email"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "access_control_role": {
        "name": "admin"
      },
      "certificate": {
        "ca": "VeriSign",
        "name": "string",
        "serial_number": 1234567890
      },
      "connectivity": {
        "errors": [
          {
            "message": {
              "arguments": [
                {
                  "code": "string",
                  "message": "string"
                }
              ],
              "code": "4",
              "message": "entry doesn't exist"
            },
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ],
        "state": "fail"
      },
      "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
      "filters": [
        {
          "name": "important-events"
        }
      ],
      "name": "Admin_Email",
      "syslog": {
        "format": {
          "hostname_override": "string",
```

```
            "message": "string",
            "timestamp_override": "string"
          },
          "port": 514,
          "transport": "string"
        },
        "system_defined": 1,
        "type": "email"
      }
    ]
  }
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983104 | The syslog destination provided is invalid |
| 983107 | A provided filter does not exist |
| 983116 | The number of notifications has reached maximum capacity |
| 983117 | The number of destinations has reached maximum capacity |
| 983129 | The rest-api destination provided must contain a valid scheme, such as http// or https// |
| 983130 | The rest-api destination provided contains an invalid URL |
| 983131 | The rest-api destination provided contains an invalid IPv6 URL |
| 983144 | The security certificate information provided is incomplete. Provide the certificate and serial number |
| 983145 | The rest-api destination provided has an 'http://' scheme. It is invalid to provide certificate information |

| Error Code | Description |
| --- | --- |
| 983149 | New SNMP destinations cannot be created |
| 983152 | Default destinations cannot be modified or removed |
| 983153 | The security certificate provided does not exist |
| 983154 | The necessary private key is not installed on the system |
| 983170 | Connectivity check is not supported for the specified destination type on the node |
| 983171 | Connectivity check failed for the specified destination on the node |
| 983176 | Connectivity check is only supported for TCP-based syslog destinations |
| 983179 | The value for the destination field cannot be empty |
| 983180 | The destination name provided cannot be empty |
| 983181 | The destination name provided cannot contain spaces |
| 983182 | The destination name provided is invalid. The destination name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or _(underscore). The allowed special characters are _(underscore) and -(hyphen) |
| 983184 | A provided property cannot be configured on the requested destination type |
| 983200 | Access control role compatibility issue with provided filters |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

certificate

Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ca | string | Client certificate issuing CA |
| name | string | Certificate name |
| serial_number | string | Client certificate serial number |

arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

message

Information to be displayed to the user.

| Name | Type | Description |
|---|---|---|
| arguments | array[arguments] | Message arguments |
| code | string | Unique message code. |
| message | string | User message. |

node

| Name | Type | Description |
|---|---|---|
| name | string | |
| uuid | string | |

errors

Error object included in the event of connectivity failure.

| Name | Type | Description |
|---|---|---|
| message | message | Information to be displayed to the user. |
| node | node | |

connectivity

| Name | Type | Description |
|---|---|---|
| errors | array[errors] | A list of errors encountered during connectivity checks. |
| state | string | Current connectivity state. |

filters

| Name | Type | Description |
|---|---|---|
| name | string | |

format

| Name | Type | Description |
| --- | --- | --- |
| hostname_override | string | Syslog Hostname Format Override. The supported hostname formats are no_override (hostname format based on the syslog.format.message property i.e. fqdn if syslog.format.message is rfc_5424, hostname_only if syslog.format.message is legacy_netapp), fqdn (Fully Qualified Domain Name) and hostname_only. |
| message | string | Syslog Message Format. The supported message formats are legacy_netapp (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG) and rfc_5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG). |
| timestamp_override | string | Syslog Timestamp Format Override. The supported timestamp formats are no_override (timestamp format based on the syslog.format.message property i.e. rfc_3164 if syslog.format.message is legacy_netapp, iso_8601_local_time if syslog.format.message is rfc_5424), rfc_3164 (format: Mmm dd hh:mm:ss), iso_8601_local_time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm) and iso_8601_utc (format: YYYY-MM-DDThh:mm:ssZ). |

syslog

| Name | Type | Description |
| --- | --- | --- |
| format | format | |
| port | integer | Syslog Port. |

| Name | Type | Description |
| --- | --- | --- |
| transport | string | Syslog Transport Protocol. |

ems_destination

| Name | Type | Description |
| --- | --- | --- |
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| name | string | Destination name. Valid in POST. |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |
| type | string | Type of destination. Valid in POST. |

records

| Name | Type | Description |
| --- | --- | --- |
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| name | string | Destination name. Valid in POST. |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |
| type | string | Type of destination. Valid in POST. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage an event destination instance

### Manage an event destination instance

#### Overview

Manages a specific instance of a destination. There are limits to the information that you can modify after a destination is created. For example, you cannot change a destination's type, but you can modify the underlying details of the type.

> The system defines default destinations that cannot be removed or modified. These destinations are specified by setting the "system_defined" field to "true".

See the documentation for /support/ems/destinations for details on the various properties in a destination.

#### Connectivity Test

Optionally, you can request the connectivity object by specifically requesting the object in the fields query. This will perform an additional test to determine the state of the destination. The state response can include one of the following values:

- success
- fail
- not_supported

### success

If the connectivity object indicates a state of 'success', then the destination check passed.

> Currently, only the 'rest_api' destination type is supported. A successful result indicates that the server received the event.

### fail

If the connectivity object indicates a state of 'fail', then the destination check has not passed. The object will contain a 'error' object with additional information regarding the failure.

#### not_supported

If the connectivity object indicates a state of 'not_supported', then the destination check is not available for the indicated destination type. This is not considered a failure condition.

> Currently, only the 'rest_api' destination type is supported for connectivity testing.

#### Examples

### Retrieving a specific destination instance

```
# The API:
GET /api/support/ems/destinations/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/destinations/snmp-traphost"
-H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"name": "snmp-traphost",
"type": "snmp",
"destination": "",
"filters": [
  {
    "name": "default-trap-events",
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/default-trap-events"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/ems/destinations/snmp-traphost"
  }
}
}
```

### Check whether a destination passes connectivity tests

```
# The API:
GET /api/support/ems/destinations/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/destinations/rest-api-
destination?fields=name,connectivity.*" -H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"name": "rest-api-destination",
"connectivity": {
  "state": "fail",
  "errors": [
    {
      "message": {
        "message": "Cannot reach host mail@mail.com.",
        "code": "4",
        "arguments": [
          {
            "code": "5",
            "message": "mail@mail.com"
          }
        ],
      },
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/support/ems/destinations/rest-api-
destination?fields=name,connectivity.*"
  }
}
}
```

**Updating an existing destination (change of email address)**

```
# The API:
PATCH /api/support/ems/destinations/{name}

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/destinations/test-
destination" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "@test_ems_destinations_patch.txt"
test_ems_destinations_patch.txt(body):
{
"destination": "support@mycompany.com"
}

# The response:
200 OK
```

#### Deleting an existing destination

```
# The API:
DELETE /api/support/ems/destinations/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/ems/destinations/test-
destination" -H "accept: application/hal+json"

# The response:
200 OK
```

### Delete an event destination

DELETE `/support/ems/destinations/{name}`

**Introduced In:** 9.6

Deletes an event destination.

#### Related ONTAP commands

- `event notification destination delete`
- `event notification delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Destination name |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983152 | Default destinations cannot be modified or removed |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

#### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an event destination

GET `/support/ems/destinations/{name}`

**Introduced In:** 9.6

Retrieves event destinations.

#### Expensive properties

There is an added computational cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter.

- `connectivity.*`

#### Related ONTAP commands

- `event notification destination show`
- `event notification show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Destination name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| name | string | Destination name. Valid in POST. |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |
| type | string | Type of destination. Valid in POST. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_control_role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ca": "VeriSign",
    "name": "string",
    "serial_number": 1234567890
  },
  "connectivity": {
    "errors": [
      {
        "message": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
```

```
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "state": "fail"
    },
    "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
    "filters": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "important-events"
      }
    ],
    "name": "Admin_Email",
    "syslog": {
      "format": {
        "hostname_override": "string",
        "message": "string",
        "timestamp_override": "string"
      },
      "port": 514,
      "transport": "string"
    },
    "system_defined": 1,
    "type": "email"
  }
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

certificate

Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ca | string | Client certificate issuing CA |
| name | string | Certificate name |
| serial_number | string | Client certificate serial number |

arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

message

Information to be displayed to the user.

| Name | Type | Description |
|---|---|---|
| arguments | array[arguments] | Message arguments |
| code | string | Unique message code. |
| message | string | User message. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |
| uuid | string | |

errors

Error object included in the event of connectivity failure.

| Name | Type | Description |
|---|---|---|
| message | message | Information to be displayed to the user. |
| node | node | |

connectivity

| Name | Type | Description |
|---|---|---|
| errors | array[errors] | A list of errors encountered during connectivity checks. |
| state | string | Current connectivity state. |

filters

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

format

| Name | Type | Description |
| --- | --- | --- |
| hostname_override | string | Syslog Hostname Format Override. The supported hostname formats are no_override (hostname format based on the syslog.format.message property i.e. fqdn if syslog.format.message is rfc_5424, hostname_only if syslog.format.message is legacy_netapp), fqdn (Fully Qualified Domain Name) and hostname_only. |
| message | string | Syslog Message Format. The supported message formats are legacy_netapp (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG) and rfc_5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG). |
| timestamp_override | string | Syslog Timestamp Format Override. The supported timestamp formats are no_override (timestamp format based on the syslog.format.message property i.e. rfc_3164 if syslog.format.message is legacy_netapp, iso_8601_local_time if syslog.format.message is rfc_5424), rfc_3164 (format: Mmm dd hh:mm:ss), iso_8601_local_time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm) and iso_8601_utc (format: YYYY-MM-DDThh:mm:ssZ). |

syslog

| Name | Type | Description |
| --- | --- | --- |
| format | format | |
| port | integer | Syslog Port. |

| Name | Type | Description |
| --- | --- | --- |
| transport | string | Syslog Transport Protocol. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an event destination

PATCH `/support/ems/destinations/{name}`

**Introduced In:** 9.6

Updates an event destination.

#### Recommended optional properties

- `filters.name` - New list of filters that should direct to this destination. The existing list is discarded.
- `certificate` - New certificate parameters when the destination type is `rest api`.

#### Related ONTAP commands

- `event notification destination modify`
- `event notification modify`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Destination name |

**Request Body**

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |

**Example request**

```
{
  "access_control_role": {
    "name": "admin"
  },
  "certificate": {
    "ca": "VeriSign",
    "name": "string",
    "serial_number": 1234567890
  },
  "connectivity": {
    "errors": [
      {
        "message": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "state": "fail"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ],
  "syslog": {
    "format": {
      "hostname_override": "string",
      "message": "string",
      "timestamp_override": "string"
    },
    "port": 514,
```

```
    "transport": "string"
  },
  "system_defined": 1
}
```

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |

**Example response**


```
{
  "access_control_role": {
    "name": "admin"
  },
  "certificate": {
    "ca": "VeriSign",
    "name": "string",
    "serial_number": 1234567890
  },
  "connectivity": {
    "errors": [
      {
        "message": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        },
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "state": "fail"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ],
  "syslog": {
    "format": {
      "hostname_override": "string",
      "message": "string",
      "timestamp_override": "string"
    },
    "port": 514,
```

```
    "transport": "string"
  },
  "system_defined": 1
}
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983088 | The destination name provided cannot be empty |
| 983089 | The destination name provided cannot contain spaces |
| 983094 | The destination name provided is invalid. The destination name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or _(underscore). The allowed special characters are _(underscore) and -(hyphen) |
| 983104 | The syslog destination provided is invalid |
| 983107 | A provided filter does not exist |
| 983116 | The number of notifications has reached maximum capacity |
| 983117 | The number of destinations has reached maximum capacity |
| 983129 | The rest-api destination must contain a valid scheme, such as http// or https// |
| 983130 | The rest-api destination provided contains an invalid URL |
| 983131 | The rest-api destination provided contains an invalid IPv6 URL |
| 983142 | The security certificate provided does not exist |
| 983143 | The private security key provided does not exist |
| 983144 | The security certificate information provided is incomplete. Provide the certificate and serial number |
| 983145 | The rest-api destination provided has an 'http://' scheme. It is invalid to provide certificate information |
| 983150 | The type of an existing destination cannot be changed |
| 983152 | Default destinations cannot be modified or removed |

| Error Code | Description |
| --- | --- |
| 983184 | A provided property cannot be configured on the requested destination type |
| 983200 | Access control role compatibility issue with provided filters |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

certificate

Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| ca | string | Client certificate issuing CA |
| name | string | Certificate name |
| serial_number | string | Client certificate serial number |

arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

message

Information to be displayed to the user.

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[arguments] | Message arguments |
| code | string | Unique message code. |
| message | string | User message. |

node

| Name | Type | Description |
| --- | --- | --- |
| name | string | |
| uuid | string | |

errors

Error object included in the event of connectivity failure.

| Name | Type | Description |
| --- | --- | --- |
| message | message | Information to be displayed to the user. |
| node | node | |

connectivity

| Name | Type | Description |
| --- | --- | --- |
| errors | array[errors] | A list of errors encountered during connectivity checks. |
| state | string | Current connectivity state. |

filters

| Name | Type | Description |
| --- | --- | --- |
| name | string | |

format

| Name | Type | Description |
| --- | --- | --- |
| hostname_override | string | Syslog Hostname Format Override. The supported hostname formats are no_override (hostname format based on the syslog.format.message property i.e. fqdn if syslog.format.message is rfc_5424, hostname_only if syslog.format.message is legacy_netapp), fqdn (Fully Qualified Domain Name) and hostname_only. |
| message | string | Syslog Message Format. The supported message formats are legacy_netapp (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG) and rfc_5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG). |
| timestamp_override | string | Syslog Timestamp Format Override. The supported timestamp formats are no_override (timestamp format based on the syslog.format.message property i.e. rfc_3164 if syslog.format.message is legacy_netapp, iso_8601_local_time if syslog.format.message is rfc_5424), rfc_3164 (format: Mmm dd hh:mm:ss), iso_8601_local_time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm) and iso_8601_utc (format: YYYY-MM-DDThh:mm:ssZ). |

syslog

| Name | Type | Description |
| --- | --- | --- |
| format | format | |
| port | integer | Syslog Port. |

| Name | Type | Description |
| --- | --- | --- |
| transport | string | Syslog Transport Protocol. |

ems_destination

| Name | Type | Description |
| --- | --- | --- |
| access_control_role | access_control_role | Indicates the access control role that created the event destination and is used to control access to the destination based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| certificate | certificate | Specifies the client-side certificate used by the ONTAP system when mutual authentication is required. This object is only applicable for **rest_api** type destinations. Both the `ca` and `serial_number` fields must be specified when configuring a certificate in a PATCH or POST request. The `name` field is read-only and cannot be used to configure a client-side certificate.<br>• Introduced in: 9.6 |
| connectivity | connectivity | |
| destination | string | Event destination |
| filters | array[filters] | |
| syslog | syslog | |
| system_defined | boolean | Flag indicating system-defined destinations. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve events

### Retrieve events

#### Overview

Queries a live collection of observed events on the system.
Note: The `filter.name` parameter is used to pass a filter to be applied to the event collection returned. To retrieve the list of available filters, refer to /support/ems/filters

#### Example

### Querying for the latest event received by EMS

```
# The API:
GET /api/support/ems/events

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/ems/events?fields=message.name&max_records=1" -H "accept:
application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "node": {
      "name": "node1",
      "uuid": "f087b8e3-99ac-11e8-b5a5-005056bb4ec7",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/f087b8e3-99ac-11e8-b5a5-
005056bb4ec7"
        }
      }
    },
    "index": 661,
    "message": {
      "name": "raid.aggr.log.CP.count"
    },
    "_links": {
      "self": {
        "href": "/api/support/ems/events/node1/661"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/support/ems/events?fields=message.name&max_records=1"
  },
}
}
```

### Retrieve observed events

GET `/support/ems/events`

**Introduced In:** 9.6

Retrieves a collection of observed events.

> The default behavior is to filter 'DEBUG' severity events. If those events are required, a severity query must be provided.

#### Related ONTAP commands

- `event log show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| time | string | query | False | Filter by time |
| index | integer | query | False | Filter by index |
| message.severity | string | query | False | Filter by message.severity |
| message.name | string | query | False | Filter by message.name |
| parameters.name | string | query | False | Filter by parameters.name |
| parameters.value | string | query | False | Filter by parameters.value |
| source | string | query | False | Filter by source |
| log_message | string | query | False | Filter by log_message |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| filter.name | string | query | False | Filter the collection returned using an event filter<br><br>• Introduced in: 9.8 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br><br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": 1,
      "log_message": "string",
      "message": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "callhome.spares.low",
        "severity": "emergency"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "parameters": [
        {
          "name": "numOps",
          "value": 123
        }
      ],
      "source": "string",
```

```
      "time": "string"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983093 | The provided filter is unknown. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

self

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | self | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

message

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Message name of the event. Returned by default. |
| severity | string | Severity of the event. Returned by default. |

node

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

| Name | Type | Description |
| --- | --- | --- |
| uuid | string | |

parameters

| Name | Type | Description |
| --- | --- | --- |
| name | string | Name of parameter |
| value | string | Value of parameter |

records

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| index | integer | Index of the event. Returned by default. |
| log_message | string | A formatted text string populated with parameter details. Returned by default. |
| message | message | |
| node | node | |
| parameters | array[parameters] | A list of parameters provided with the EMS event. |
| source | string | Source |
| time | string | Timestamp of the event. Returned by default. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
| --- | --- | --- |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Create or retrieve event filters

### Create or retrieve event filters

#### Overview

Manages the list of available filters. A filter is a named collection of rules that enable the system to identify events that require additional handling. A filter is linked with a destination to which the system sends specific events.

When EMS processes an event, each filter is evaluated for a match. More than one filter can handle a single event.

> The system defines default filters that cannot be removed or modified. These filters are specified by setting the "system_defined" field to "true".

#### Filter rule position

A filter's rules are evaluated sequentially, according to their position index. When a rule is added or modified, the position can be set to customize the filter's logic. If no position is specified, a new rule is appended to the end of the list.

#### Filter rule types

A filter rule can be one of two types: 'include' or 'exclude'.
If an event matches the criteria of the rule, the type dictates whether it should be forwarded to the destination or ignored.

#### Filter rule matching criteria

A valid filter rule must contain at least one set of criteria.

#### Name pattern

A name pattern is matched against an event's name. Multiple characters can be matched using the wildcard character '*'.

#### Severity

The severity pattern is matched against an event's severity. Multiple severities can be specified in a comma separated list. A single wildcard `*` will match all severities.
When multiple severities are provided in a rule, all must match for the rule to be considered matched. A pattern

can include one or more wildcard * characters. Valid values are:

- emergency
- alert
- error
- notice
- informational
- debug

### SNMP trap type

The SNMP trap type pattern is matched against an event's trap type. Multiple trap types can be specified in a comma separated list. A single wildcard * matches all trap types.
When multiple trap types are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard * characters. Valid values are:

- standard
- built_in
- severity_based

### Parameter criteria

A parameter criterion is matched against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters.

#### Examples

### Retrieving a list of filters whose names contain a hyphen

```
# The API:
GET /api/support/ems/filters

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/filters?name=*-*" -H
"accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "name": "default-trap-events",
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/default-trap-events"
      }
    }
  },
  {
    "name": "important-events",
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/important-events"
      }
    }
  },
  {
    "name": "no-info-debug-events",
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/no-info-debug-events"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/support/ems/filters?name=*-*"
  }
}
}
```

**Creating a new filter using various matching criteria**

```
# The API:
POST /api/support/ems/filters

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/filters" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_filters_post.txt"
test_ems_filters_post.txt(body):
{
"name": "test-filter",
"rules": [
  {
    "index": 1,
    "type": "include",
    "message_criteria": {
      "name_pattern": "LUN.*",
      "severities": "alert,error",
      "snmp_trap_types": "severity_based"
    },
    "parameter_criteria": [
      {
        "name_pattern": "type",
        "value_pattern": "volume"
      },
      {
        "name_pattern": "vol",
        "value_pattern": "cloud*"
      }
    ]
  }
]
}

# The response:
201 Created
```

**Retrieve event filters**

GET `/support/ems/filters`

**Introduced In:** 9.6

Retrieves a collection of event filters.

#### Related ONTAP commands

- `event filter show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| system_defined | boolean | query | False | Filter by system_defined<br>• Introduced in: 9.10 |
| access_control_role.name | string | query | False | Filter by access_control_role.name<br>• Introduced in: 9.13 |
| rules.message_criteria.snmp_trap_types | string | query | False | Filter by rules.message_criteria.snmp_trap_types |
| rules.message_criteria.severities | string | query | False | Filter by rules.message_criteria.severities |
| rules.message_criteria.name_pattern | string | query | False | Filter by rules.message_criteria.name_pattern |
| rules.parameter_criteria.name_pattern | string | query | False | Filter by rules.parameter_criteria.name_pattern<br>• Introduced in: 9.13 |
| rules.parameter_criteria.value_pattern | string | query | False | Filter by rules.parameter_criteria.value_pattern<br>• Introduced in: 9.13 |
| rules.type | string | query | False | Filter by rules.type |
| rules.index | integer | query | False | Filter by rules.index |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | query | False | Filter by name |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_control_role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "name": "wafl-critical-events",
      "rules": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "index": 1,
          "message_criteria": {
            "_links": {
              "related": {
                "href": "/api/resourcelink"
              }
            },
            "name_pattern": "wafl.*",
            "severities": "emergency,alert,error",
            "snmp_trap_types": "standard,built_in"
          },
```

```
          "parameter_criteria": [
            {
              "name_pattern": "vol",
              "value_pattern": "cloud*"
            }
          ],
          "type": "include"
        }
      ],
      "system_defined": 1
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

_links

| Name | Type | Description |
|---|---|---|
| related | href | |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |

| Name | Type | Description |
| --- | --- | --- |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
| --- | --- | --- |
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

rules

Rule for an event filter

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| name | string | Filter name |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Create an event filter**

POST `/support/ems/filters`

**Introduced In:** 9.6

Creates an event filter.

#### Required properties

- `name` - String that uniquely identifies the filter.

#### Recommended optional properties

- `rules` - List of criteria which is used to match a filter with an event.

#### Related ONTAP commands

- `event filter create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| name | string | Filter name |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

### Example request

```
{
  "access_control_role": {
    "name": "admin"
  },
  "name": "wafl-critical-events",
  "rules": [
    {
      "index": 1,
      "message_criteria": {
        "name_pattern": "wafl.*",
        "severities": "emergency,alert,error",
        "snmp_trap_types": "standard,built_in"
      },
      "parameter_criteria": [
        {
          "name_pattern": "vol",
          "value_pattern": "cloud*"
        }
      ],
      "type": "include"
    }
  ],
  "system_defined": 1
}
```

### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "num_records": 1,
  "records": [
    {
      "access_control_role": {
        "name": "admin"
      },
      "name": "wafl-critical-events",
      "rules": [
        {
          "index": 1,
          "message_criteria": {
            "name_pattern": "wafl.*",
            "severities": "emergency,alert,error",
            "snmp_trap_types": "standard,built_in"
          },
          "parameter_criteria": [
            {
              "name_pattern": "vol",
              "value_pattern": "cloud*"
            }
          ],
          "type": "include"
        }
      ],
      "system_defined": 1
    }
  ]
}
```


### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 983088 | The filter name provided is empty |
| 983089 | The filter name provided cannot contain spaces |
| 983092 | The index of the rule provided is outside the allowed range for the filter provided |
| 983094 | The filter name provided is invalid. The filter name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or (underscore). The allowed special characters are (underscore) and -(hyphen) |
| 983095 | The rule index provided is invalid for the filter provided |
| 983101 | No event is matched by the rule provided |
| 983113 | Default filters cannot be modified or removed |
| 983114 | The maximum number of filters is reached |
| 983115 | The maximum number of filter rules is reached |
| 983126 | A rule requires at least one name_pattern, severities, snmp_trap_types, or parameter pattern to be defined |
| 983127 | A property cannot contain a combination of the wildcard character and other values |
| 983128 | An invalid value is provided for the property 'snmp_trap_types' |
| 983146 | An invalid value is provided for the property 'severities' |
| 983147 | The severities provided are not supported |
| 983155 | The provided severities property does not match that of the name_pattern |
| 983156 | The provided snmp_trap_types property does not match that of the name_pattern |
| 983157 | The provided severities and snmp_trap_types properties do not match those of the name_pattern |
| 983158 | The name_pattern provided does not exist |
| 983195 | Empty field in parameter_criteria. Both name and value patterns must be specified |
| 983196 | name_pattern and value_pattern fields in parameter_criteria are empty |
| 983211 | Parameter criteria based filtering is not supported in this version of ONTAP |
| 983216 | The provided parameter criteria does not match that of the message name |

Also see the table of common errors in the Response body overview section of this documentation.

## Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

rules

Rule for an event filter

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

ems_filter

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| name | string | Filter name |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

records

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| name | string | Filter name |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage event filters

### Manage event filters

#### Overview

Manages a specific filter instance.
See the documentation for /support/ems/filters for details on the various properties.

#### Examples

### Retrieving a specific filter instance

```
# The API:
GET /api/support/ems/filters/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/filters/aggregate-events"
-H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"name": "aggregate-events",
"rules": [
  {
    "index": 1,
    "type": "include",
    "message_criteria": {
      "name_pattern": "*",
      "severities": "emergency,alert,error,notice",
      "snmp_trap_types": "*",
      "_links": {
        "related": {
          "href":
"/api/support/ems/messages?name=*&severity=emergency,alert,error,notice&sn
mp_trap_type=*"
        }
      }
    },
    "parameter_criteria": [
      {
        "name_pattern": "type",
        "value_pattern": "aggregate"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/aggregate-events/rules/1"
      }
    }
  },
  {
```

```
    "index": 2,
    "type": "exclude",
    "message_criteria": {
      "name_pattern": "*",
      "severities": "*",
      "snmp_trap_types": "*",
      "_links": {
        "related": {
          "href":
"/api/support/ems/messages?name=*&severity=*&snmp_trap_type=*"
        }
      }
    },
    "parameter_criteria": [
      {
        "name_pattern": "*",
        "value_pattern": "*"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/aggregate-events/rules/2"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/ems/filters/aggregate-events"
  }
}
}
```


**Updating an existing filter with a new rule**

```
# The API:
PATCH /api/support/ems/filters/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/ems/filters/test-filter" -H
"accept: application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_filters_patch.txt"
test_ems_filters_patch.txt(body):
{
"rules": [
  {
    "type": "include",
    "message_criteria": {
      "name_pattern": "wafl.*",
      "severities": "error"
    }
  }
]
}

# The response:
200 OK
```

**Deleting an existing filter**

```
# The API:
DELETE /api/support/ems/filters/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/ems/filters/test-filter" -H
"accept: application/hal+json"

# The response:
200 OK
```

**Delete an event filter**

DELETE `/support/ems/filters/{name}`

**Introduced In:** 9.6

Deletes an event filter.

#### Related ONTAP commands

- `event filter delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983113 | Default filters cannot be modified or removed |
| 983124 | Filter is being referenced by a destination |
| 983204 | Filter is being used for role-based operations |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an event filter

GET `/support/ems/filters/{name}`

**Introduced In:** 9.6

Retrieves an event filter.

#### Related ONTAP commands

- `event filter show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| name | string | Filter name |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_control_role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "name": "wafl-critical-events",
  "rules": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": 1,
      "message_criteria": {
        "_links": {
          "related": {
            "href": "/api/resourcelink"
          }
        },
        "name_pattern": "wafl.*",
        "severities": "emergency,alert,error",
        "snmp_trap_types": "standard,built_in"
      },
      "parameter_criteria": [
        {
          "name_pattern": "vol",
          "value_pattern": "cloud*"
        }
      ],
      "type": "include"
    }
  ],
  "system_defined": 1
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

_links

| Name | Type | Description |
|---|---|---|
| related | href | |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
| --- | --- | --- |
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

rules

Rule for an event filter

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an event filter

PATCH `/support/ems/filters/{name}`

**Introduced In:** 9.6

Updates an event filter.

#### Recommended optional properties

- `new_name` - New string that uniquely identifies a filter.
- `rules` - New list of criteria used to match the filter with an event. The existing list is discarded.

#### Related ONTAP commands

- `event filter rename`
- `event filter rule add`
- `event filter rule delete`
- `event filter rule reorder`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |
| new_name | string | query | False | New filter name for renames. Valid in PATCH. |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

**Example request**

```
{
  "access_control_role": {
    "name": "admin"
  },
  "rules": [
    {
      "index": 1,
      "message_criteria": {
        "name_pattern": "wafl.*",
        "severities": "emergency,alert,error",
        "snmp_trap_types": "standard,built_in"
      },
      "parameter_criteria": [
        {
          "name_pattern": "vol",
          "value_pattern": "cloud*"
        }
      ],
      "type": "include"
    }
  ],
  "system_defined": 1
}
```

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983088 | The filter name provided is empty |
| 983089 | The filter name provided cannot contain spaces |
| 983092 | The index of the rule provided is outside the allowed range for the filter provided |
| 983094 | The filter name provided is invalid. The filter name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or (underscore). The allowed special characters are (underscore) and -(hyphen) |
| 983095 | The rule index provided is invalid for the filter provided |
| 983101 | No event is matched by the rule provided |
| 983113 | Default filters cannot be modified or removed |
| 983114 | The maximum number of filters is reached |
| 983115 | The maximum number of filter rules is reached |
| 983126 | A rule requires at least one name_pattern, severities, snmp_trap_types, or parameter pattern to be defined |
| 983127 | A property cannot contain a combination of the wildcard characters and other values |
| 983128 | An invalid value is provided for the property 'snmp_trap_types' |
| 983146 | An invalid value is provided for the property 'severities' |
| 983147 | The severities provided are not supported |
| 983155 | The provided severities property does not match that of the name_pattern |
| 983156 | The provided snmp_trap_types property does not match that of the name_pattern |
| 983157 | The provided severities and snmp_trap_types properties do not match those of the name_pattern |
| 983158 | The name_pattern provided does not exist |
| 983195 | Empty field in parameter_criteria. Both name and value patterns must be specified |
| 983196 | name_pattern and value_pattern fields in parameter_criteria are empty |

| Error Code | Description |
| --- | --- |
| 983211 | Parameter criteria based filtering is not supported in this version of ONTAP |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

rules

Rule for an event filter

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

ems_filter

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Indicates the access control role that created the event filter and is used to control access to the filter based on role-based access control (RBAC) rules. If created by the 'admin' user, the field is unset. |
| rules | array[rules] | Array of event filter rules on which to match. |
| system_defined | boolean | Flag indicating system-defined filters. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Create or retrieve filter rules

### Create or retrieve filter rules

#### Overview

Manages the list of rules associated with a specific filter. A filter contains a list of rules that are evaluated to determine whether an event matches the filter. When a rule matches an event, the filter is considered a match. See the documentation for /support/ems/filters for details on the various properties in a rule.

#### Examples

### Retrieving the collection of rules connected to a filter

```
# The API:
GET /api/support/ems/filters/{name}/rules

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/filters/no-info-debug-
events/rules" -H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "index": 1,
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/no-info-debug-events/rules/1"
      }
    }
  },
  {
    "index": 2,
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/no-info-debug-events/rules/2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/support/ems/filters/no-info-debug-events/rules"
  }
}
}
```

**Creating a new rule under an existing filter**

The absence of a position index appends the rule to the end of the list.

```
# The API:
POST /api/support/ems/filters/{name}/rules

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/filters/test-filter/rules"
-H "accept: application/hal+json" -H "Content-Type: application/json" -d
"@test_ems_filter_rules_post.txt"
test_ems_filter_rules_post.txt(body):

# JSON Body
{
"type": "include",
"message_criteria": {
  "name_pattern": "wafl.*",
  "severities": "error,informational"
}
}

# The response:
201 Created
```

**Retrieve event filter rules**

GET `/support/ems/filters/{name}/rules`

**Introduced In:** 9.6

Retrieves event filter rules.

#### Related ONTAP commands

- `event filter show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter Name |
| message_criteria.snmp_trap_types | string | query | False | Filter by message_criteria.snmp_trap_types |
| message_criteria.severities | string | query | False | Filter by message_criteria.severities |

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| message_criteria.name_pattern | string | query | False | Filter by message_criteria.name_pattern |
| parameter_criteria.name_pattern | string | query | False | Filter by parameter_criteria.name_pattern<br><br>• Introduced in: 9.13 |
| parameter_criteria.value_pattern | string | query | False | Filter by parameter_criteria.value_pattern<br><br>• Introduced in: 9.13 |
| type | string | query | False | Filter by type |
| index | integer | query | False | Filter by index |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br><br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

### Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": 1,
      "message_criteria": {
        "_links": {
          "related": {
            "href": "/api/resourcelink"
          }
        },
        "name_pattern": "wafl.*",
        "severities": "emergency,alert,error",
        "snmp_trap_types": "standard,built_in"
      },
      "parameter_criteria": [
        {
          "name_pattern": "vol",
          "value_pattern": "cloud*"
        }
      ],
      "type": "include"
    }
  ]
}
```


### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| related | href | |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
| --- | --- | --- |
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

records

Rule for an event filter

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an event filter rule

POST `/support/ems/filters/{name}/rules`

**Introduced In:** 9.6

Creates an event filter rule.

#### Required properties

- `type` - Enumeration indicating whether the rule is for include or exclude.
- `message_criteria` - Message criteria used by a rule to match an event.
- `parameter_criteria` - Parameter criteria used by a rule to match an event. Note: At least one pattern needs to be provided for message_criteria and / or parameter_criteria.

#### Recommended optional properties

- `index` - One-based position index of the new rule.

#### Related ONTAP commands

- `event filter rule add`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

**Request Body**

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

**Example request**

```
{
  "index": 1,
  "message_criteria": {
    "name_pattern": "wafl.*",
    "severities": "emergency,alert,error",
    "snmp_trap_types": "standard,built_in"
  },
  "parameter_criteria": [
    {
      "name_pattern": "vol",
      "value_pattern": "cloud*"
    }
  ],
  "type": "include"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "num_records": 1,
  "records": [
    {
      "index": 1,
      "message_criteria": {
        "name_pattern": "wafl.*",
        "severities": "emergency,alert,error",
        "snmp_trap_types": "standard,built_in"
      },
      "parameter_criteria": [
        {
          "name_pattern": "vol",
          "value_pattern": "cloud*"
        }
      ],
      "type": "include"
    }
  ]
}
```


**Headers**

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 983092 | The index of the rule provided is outside the allowed range for the filter provided |
| 983095 | The rule index provided is invalid for the filter provided |
| 983113 | Default filters cannot be modified or removed |
| 983115 | The maximum number of filter rules is reached |
| 983126 | A rule requires at least one name_pattern, severities, snmp_trap_types, or parameter pattern to be defined |
| 983127 | A property cannot contain a combination of the wildcard characters and other values |
| 983128 | An invalid value is provided for the property 'snmp_trap_types' |
| 983146 | An invalid value is provided for the property 'severities' |
| 983147 | The severities provided are not supported |
| 983155 | The provided severities property does not match that of the name_pattern |
| 983156 | The provided snmp_trap_types property does not match that of the name_pattern |
| 983157 | The provided severities and snmp_trap_types properties do not match those of the name_pattern |
| 983158 | The name_pattern provided does not exist |
| 983195 | Empty field in parameter_criteria. Both name and value patterns must be specified |
| 983196 | name_pattern and value_pattern fields in parameter_criteria are empty |
| 983211 | Parameter criteria based filtering is not supported in this version of ONTAP |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

ems_filter_rule

Rule for an event filter

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |

| Name | Type | Description |
| --- | --- | --- |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

records

Rule for an event filter

| Name | Type | Description |
| --- | --- | --- |
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage event filter rules

### Manage event filter rules

#### Overview

Manages a specific instance of a rule within a filter.
See the documentation for /support/ems/filters for details on the various properties in a rule.

#### Examples

### Retrieving a single instance of a rule

```
# The API:
GET /api/support/ems/filters/{name}/rules/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/filters/no-info-debug-
events/rules/1" -H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"name": "no-info-debug-events",
"index": 1,
"type": "include",
"message_criteria": {
  "name_pattern": "*",
  "severities": "emergency,alert,error,notice",
  "snmp_trap_types": "*",
  "_links": {
    "self": {
      "href":
"/api/support/ems/messages?name=*&severity=emergency,alert,error,notice&sn
mp_trap_type=*"
    }
  }
},
"parameter_criteria": [
  {
    "name_pattern": "*",
    "value_pattern": "*"
  }
],
"_links": {
  "self": {
    "href": "/api/support/ems/filters/no-info-debug-events/rules/1"
  }
}
}
```

**Updating an existing rule to use severity emergency**

```
# The API:
PATCH /api/support/ems/filters/{name}/rules/{index}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/ems/filters/test-
filter/rules/1" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "@test_ems_filter_rules_patch.txt"
test_ems_filter_rules_patch.txt(body):
{
"message_criteria": {
  "severities": "emergency"
}
}

# The response:
200 OK
```

#### Deleting a rule from an existing filter

```
# The API:
DELETE /api/support/ems/filters/{name}/rules/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/ems/filters/test-
filter/rules/1" -H "accept: application/hal+json"

# The response:
200 OK
```

### Delete an event filter rule

DELETE `/support/ems/filters/{name}/rules/{index}`

**Introduced In:** 9.6

Deletes an event filter rule.

#### Related ONTAP commands

- `event filter rule delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| index | string | path | True | Filter rule index |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983091 | A default rule cannot be removed |
| 983092 | The index of the rule provided is outside the allowed range for the filter provided |
| 983095 | The rule index provided is invalid for the filter provided |
| 983110 | There are no user defined rules in the filter provided |
| 983113 | Default filters cannot be modified or removed |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

### See Definitions

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve an event filter rule

GET `/support/ems/filters/{name}/rules/{index}`

**Introduced In:** 9.6

Retrieves an event filter rule.

#### Related ONTAP commands

- `event filter show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |
| index | string | path | True | Filter rule index |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "index": 1,
  "message_criteria": {
    "_links": {
      "related": {
        "href": "/api/resourcelink"
      }
    },
    "name_pattern": "wafl.*",
    "severities": "emergency,alert,error",
    "snmp_trap_types": "standard,built_in"
  },
  "parameter_criteria": [
    {
      "name_pattern": "vol",
      "value_pattern": "cloud*"
    }
  ],
  "type": "include"
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| related | href | |

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update an event filter rule

PATCH `/support/ems/filters/{name}/rules/{index}`

**Introduced In:** 9.6

Updates an event filter rule.

#### Recommended optional properties

- `message_criteria` - Message criteria used by a rule to match an event.
- `parameter_criteria` - Parameter criteria used by a rule to match an event.

#### Related ONTAP commands

- `event filter rule add`
- `event filter rule delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| name | string | path | True | Filter name |
| index | string | path | True | Filter rule index |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| new_index | integer | query | False | New position for the filter rule index |

**Request Body**

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

#### Example request

```
{
  "index": 1,
  "message_criteria": {
    "name_pattern": "wafl.*",
    "severities": "emergency,alert,error",
    "snmp_trap_types": "standard,built_in"
  },
  "parameter_criteria": [
    {
      "name_pattern": "vol",
      "value_pattern": "cloud*"
    }
  ],
  "type": "include"
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983092 | The index of the rule provided is outside the allowed range for the filter provided |
| 983095 | The rule index provided is invalid for the filter provided |
| 983113 | Default filters cannot be modified or removed |
| 983126 | A rule requires at least one name_pattern, severities, snmp_trap_types, or parameter pattern to be defined |
| 983127 | A property cannot contain a combination of the wildcard characters and other values. |
| 983128 | An invalid value is provided for the property 'snmp_trap_types' |
| 983146 | An invalid value is provided for the property 'severities' |

| Error Code | Description |
| --- | --- |
| 983147 | The severities provided are not supported |
| 983155 | The provided severities property does not match that of the name_pattern |
| 983156 | The provided snmp_trap_types property does not match that of the name_pattern |
| 983157 | The provided severities and snmp_trap_types properties do not match those of the name_pattern |
| 983158 | The name_pattern provided does not exist |
| 983195 | Empty field in parameter_criteria. Both name and value patterns must be specified |
| 983196 | name_pattern and value_pattern fields in parameter_criteria are empty |
| 983211 | Parameter criteria based filtering is not supported in this version of ONTAP |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Message name filter on which to match. Supports wildcards. Defaults to * if not specified. |
| severities | string | A comma-separated list of severities or a wildcard. |
| snmp_trap_types | string | A comma separated list of snmp_trap_types or a wildcard. |

parameter_criteria

Criterion used for parameter based filtering

| Name | Type | Description |
|---|---|---|
| name_pattern | string | Parameter name pattern. Wildcard character '*' is supported. |
| value_pattern | string | Parameter value pattern. Wildcard character '*' is supported. |

ems_filter_rule

Rule for an event filter

| Name | Type | Description |
|---|---|---|
| index | integer | Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list. |

| Name | Type | Description |
|---|---|---|
| message_criteria | message_criteria | Matching message definitions for the filter. A property must be specified. |
| parameter_criteria | array[parameter_criteria] | Parameter criteria used to match against events' parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, all must match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters. |
| type | string | Rule type |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage the EMS event catalog

### Manage the EMS event catalog

#### Overview

Allows access to the EMS event catalog. The catalog contains a list of all events supported by the system and their corresponding descriptions, the reason for an event occurrence, and how to correct issues related to the event.

#### Example

### Querying for the first event that has a message name beginning with 'C'

```
# The API:
GET /api/support/ems/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/ems/messages?fields=name&max_records=1&name=C*" -H
"accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "name": "CR.Data.File.Inaccessible",
    "_links": {
      "self": {
        "href": "/api/support/ems/messages/CR.Data.File.Inaccessible"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/support/ems/messages?fields=name&max_records=1&name=C*"
  },
}
}
```

## Retrieve event catalog definitions

GET `/support/ems/messages`

**Introduced In:** 9.6

Retrieves the event catalog definitions.

#### Related ONTAP commands

- `event catalog show`

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| description | string | query | False | Filter by description |
| snmp_trap_type | string | query | False | Filter by snmp_trap_type |
| corrective_action | string | query | False | Filter by corrective_action |
| name | string | query | False | Filter by name |
| deprecated | boolean | query | False | Filter by deprecated |
| severity | string | query | False | Filter by severity |
| filter.name | string | query | False | The filter name that applies to the query.<br>• Introduced in: 9.14 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records |
| records | array[records] | |

#### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "corrective_action": "string",
      "deprecated": 1,
      "description": "string",
      "name": "callhome.spares.low",
      "severity": "error",
      "snmp_trap_type": "standard"
    }
  ]
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983093 | The provided filter is unknown. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| corrective_action | string | Corrective action |
| deprecated | boolean | Is deprecated? |
| description | string | Description of the event. |
| name | string | Name of the event. |
| severity | string | Severity |
| snmp_trap_type | string | SNMP trap type |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## View and create EMS role-based configurations

### View and create EMS role-based configurations

#### Overview

Manages the list of access control roles that have limited access to the EMS events and configurations. Access to the EMS events and configurations are controlled by the EMS role-based configuration, which has the following two properties:

- `event_filter`
- `limit_access_to_global_configs`

### Role-based event filter

A role-based event filter is a filter assigned to an access control role that limits what events users of that role have access to. It is applied transparently when a user of that role performs GET requests on the following endpoints:

- `/support/ems/events`
- `/support/ems/messages`

##
It is also applied when events are forwarded as notifications to destinations created by users of that role. To those destinations, only events that match the role-based filter in addition to the filters configured on the destinations are forwarded.

### Limiting access to the global EMS configurations

The APIs at the `/api/support/ems` endpoint expose several configuration properties that affect how EMS handles events. The following properties are always accessible by all roles that have access to the `/api/support/ems` APIs.

- `mail_from`
- `mail_server`
- `mail_server_user`

- `mail_server_password`
- `proxy_url`
- `proxy_user`
- `proxy_password`

##
Access to the following property by an access control role can be restricted:

- `pubsub_enabled`

##
Limiting access to EMS events and configurations is typically applied for an access control role that is designed to have limited administrative capabilities. For example, when the cluster is physically managed by a service provider and the services are managed by the customer.

#### Examples

**Retrieving the list of EMS role-based configurations**

```
# The API:
GET /api/support/ems/role-configs

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/role-configs" -H "accept:
application/hal+json"

# The response:
200 OK

# JSON Body
{
"records": [
  {
    "access_control_role": {
      "name": "storage-admin",
      "_links": {
        "self": {
          "href": "/api/security/roles/0b2580c8-de36-4213-bfca-
88cdaaaf3ae6/storage-admin"
        }
      }
    },
    "event_filter": {
      "name": "storage-admin-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/storage-admin-events"
```

```
        }
      }
    },
    "limit_access_to_global_configs": true,
    "_links": {
      "self": {
        "href": "/api/support/ems/role-configs/storage-admin"
      }
    }
  },
  {
    "access_control_role": {
      "name": "network-admin",
      "_links": {
        "self": {
          "href": "/api/security/roles/0b2580c8-de36-4213-bfca-
88cdaaaf3ae6/network-admin"
        }
      }
    },
    "filter_filter": {
      "name": "network-admin-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/network-admin-events"
        }
      }
    },
    "limit_access_to_global_configs": true,
    "_links": {
      "self": {
        "href": "/api/support/ems/role-configs/network-admin"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/support/ems/role-configs"
  }
}
}
```

#### Creating an EMS role-based configuration for an access control role

```
# The API:
POST /api/support/ems/role-configs

# The call:
curl -X POST "https://<mgmt-ip>/api/support/ems/role-configs" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d
"@role_configs_post.txt"
role_configs_post.txt(body):
{
"access_control_role": {
  "name": "storage-admin"
},
"event_filter": {
  "name": "storage-admin-events"
},
"limit_access_to_global_configs": true
}

# The response:
201 Created
```

### Retrieve EMS role-based configurations

GET `/support/ems/role-configs`

**Introduced In:** 9.13

Retrieves a collection of the EMS role-based configurations.

#### Related ONTAP commands

- `event role-config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| event_filter.name | string | query | False | Filter by event_filter.name |
| limit_access_to_global_configs | boolean | query | False | Filter by limit_access_to_global_configs |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| access_control_role.name | string | query | False | Filter by access_control_role.name |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[records] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_control_role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      },
      "event_filter": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "important-events"
      }
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Access control role that the event filter is assigned to.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

event_filter

Event filter that is assigned to the access control role.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

records

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Access control role that the event filter is assigned to. |
| event_filter | event_filter | Event filter that is assigned to the access control role. |

| Name | Type | Description |
|---|---|---|
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create an EMS role-based configuration for an access control role

POST `/support/ems/role-configs`

**Introduced In:** 9.13

Creates an EMS role-based configuration for an access control role.

#### Required properties

- `access_control_role` - Uniquely identifies the access control role.

#### Recommended optional properties

- `event_filter` - Identifies the event filter.
- `limit_access_to_global_configs` - Indicates whether the access control role has limited access to global EMS configurations.

#### Related ONTAP commands

- `event role-config create`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Access control role that the event filter is assigned to. |
| event_filter | event_filter | Event filter that is assigned to the access control role. |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

#### Example request

```
{
  "access_control_role": {
    "name": "admin"
  },
  "event_filter": {
    "name": "important-events"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records |
| records | array[records] | |

### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "access_control_role": {
        "name": "admin"
      },
      "event_filter": {
        "name": "important-events"
      }
    }
  ]
}
```

### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983198 | The event filter provided does not exist. |
| 983199 | The access control role provided does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Access control role that the event filter is assigned to.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

event_filter

Event filter that is assigned to the access control role.

| Name | Type | Description |
|---|---|---|
| name | string | |

ems_role_config

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Access control role that the event filter is assigned to. |
| event_filter | event_filter | Event filter that is assigned to the access control role. |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

records

| Name | Type | Description |
|---|---|---|
| access_control_role | access_control_role | Access control role that the event filter is assigned to. |
| event_filter | event_filter | Event filter that is assigned to the access control role. |

| Name | Type | Description |
| --- | --- | --- |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage EMS role-based configurations

### Manage EMS role-based configurations

#### Overview

Manages the EMS role-based configuration of a specific access control role.

##
See the documentation for /support/ems/role-configs for details on the various properties.

#### Examples

### Retrieving the EMS role-based configuration of an access control role

```
# The API:
GET /api/support/ems/role-configs/{access_control_role.name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/ems/role-configs/storage-admin"
-H "accept: application/hal+json"

# The response:
200 OK

# JSON Body
{
"access_control_role": {
  "name": "storage-admin",
  "_links": {
    "self": {
      "href": "/api/security/roles/0b2580c8-de36-4213-bfca-
88cdaaaf3ae6/storage-admin"
    }
  }
},
"event_filter": {
  "name": "storage-admin-events",
  "_links": {
    "self": {
      "href": "/api/support/ems/filters/storage-admin-events"
    }
  },
  "limit_access_to_global_configs": true
},
"_links": {
  "self": {
    "href": "/api/support/ems/role-configs/storage-admin"
  }
}
}
```

**Updating the EMS role-based configuration of an access control role**

```
# The API:
PATCH /api/support/ems/event-configs/{access_control_role.name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/ems/role-configs/storage-
admin" -H "accept: application/hal+json" -H "Content-Type:
application/json" -d "@role_configs_patch.txt"
role_configs_patch.txt(body):
{
"event_filter": {
  "name": "new-storage-admin-events"
},
"limit_access_to_global_configs": false
}

# The response:
200 OK
```

**Removing the EMS role-based configuration of an access control role**

```
# The API:
DELETE /api/support/ems/role-configs/{access_control_role.name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/ems/role-configs/storage-
admin" -H "accept: application/hal+json"

# The response:
200 OK
```

**Remove the EMS role-based configuration for an access control role**

DELETE `/support/ems/role-configs/{access_control_role.name}`

**Introduced In:** 9.13

Removes the EMS role-based configuration of the access control role.

#### Related ONTAP commands

- `event role-config delete`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| access_control_role.name | string | path | True | Access control role name |

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve the EMS role-based configuration for an access control role

GET `/support/ems/role-configs/{access_control_role.name}`

**Introduced In:** 9.13

Retrieves the EMS role-based configuration of the access control role.

#### Related ONTAP commands

- `event role-config show`

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| access_control_role.name | string | path | True | Access control role name |
| fields | array[string] | query | False | Specify the fields to return. |

#### Response

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| access_control_role | access_control_role | Access control role that the event filter is assigned to. |
| event_filter | event_filter | Event filter that is assigned to the access control role. |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

**Example response**


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_control_role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "event_filter": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "important-events"
  }
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

access_control_role

Access control role that the event filter is assigned to.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Role name |

event_filter

Event filter that is assigned to the access control role.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Update the EMS role-based configuration for an access control role

PATCH `/support/ems/role-configs/{access_control_role.name}`

**Introduced In:** 9.13

Updates the EMS role-based configuration of the access control role.

#### Recommended optional properties

- `event_filter` - Identifies the event filter.
- `limit_access_to_global_configs` - Indicates whether the access control role has limited access to global EMS configurations.

#### Related ONTAP commands

- `event role-config modify`

#### Parameters

| Name | Type | In | Required | Description |
| --- | --- | --- | --- | --- |
| access_control_role.name | string | path | True | Access control role name |

#### Request Body

| Name | Type | Description |
| --- | --- | --- |
| event_filter | event_filter | Event filter that is assigned to the access control role. |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

#### Example request

```
{
  "event_filter": {
    "name": "important-events"
  }
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 983197 | The event filter provided is the same as the currently assigned filter. |
| 983198 | The event filter provided does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

access_control_role

Access control role that the event filter is assigned to.

| Name | Type | Description |
|---|---|---|
| name | string | Role name |

event_filter

Event filter that is assigned to the access control role.

| Name | Type | Description |
|---|---|---|
| name | string | |

ems_role_config

| Name | Type | Description |
|---|---|---|
| event_filter | event_filter | Event filter that is assigned to the access control role. |
| limit_access_to_global_configs | boolean | Indicates whether the access control has limited access to global EMS configurations. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |

| Name | Type | Description |
|---|---|---|
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage cluster-wide SNMP configuration

### Manage cluster-wide SNMP configuration

#### Overview

Cluster wide SNMP configuration. You can configure or retrieve the following SNMP parameters using this endpoint:

- enable or disable SNMP
- enable or disable SNMP authentication traps
- enable or disable SNMP traps

##
This endpoint can also be used to trigger an SNMP test trap.

#### Examples

**Disables SNMP protocol in the cluster.**

```
# The API:
PATCH "/api/support/snmp"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"enabled":"false"}'

# The response
200 OK
{}
```

**Enables SNMP authentication traps in the cluster.**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d
'{"auth_traps_enabled":"true"}'

# The response
200 OK
{}
```

**Enables SNMP protocol and SNMP authentication traps in the cluster.**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"enabled":"true",
"auth_traps_enabled":"true"}'

# The response
200 OK
{}
```

**Disables the SNMP trap subsystem in the cluster. Once the SNMP trap subsystem is disabled, the cluster stops sending SNMP traps.**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"traps_enabled":"false"}'

# The response
200 OK
{}
```

**Sets the contact and location for the SNMP server**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d
'{"contact":"support@company.com", "location":"Building 1"}'

#The response
200 OK
{}
```

**Triggers an SNMP test trap.**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d
'{"trigger_test_trap":"true"}'

#The response
200 OK
{}
```

**Enables the SNMP protocol in the cluster, SNMP traps, and triggers an SNMP test trap.**

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"enabled":"true",
"traps_enabled":"true", "trigger_test_trap":"true"}'

# The response
200 OK
{}
```

### Retrieve the cluster-wide SNMP configuration

GET `/support/snmp`

**Introduced In:** 9.7

Retrieves the cluster wide SNMP configuration.

#### Related ONTAP commands

- `options snmp.enable`
- `system snmp show`

#### Learn more

- DOC /support/snmp

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | self_link | |
| auth_traps_enabled | boolean | Specifies whether to enable or disable SNMP authentication traps. |
| contact | string | Specifies the contact person for the SNMP server |

| Name | Type | Description |
|---|---|---|
| enabled | boolean | Specifies whether to enable or disable SNMP. |
| location | string | Specifies the location of the SNMP server |
| traps_enabled | boolean | Specifies whether to enable or disable SNMP traps. |
| trigger_test_trap | boolean | Trigger a test SNMP trap. |

**Example response**

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "auth_traps_enabled": 1,
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "enabled": 1,
  "location": "Building 1",
  "traps_enabled": 1,
  "trigger_test_trap": 1
}
```

**Error**

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

| Name | Type | Description |
|---|---|---|
| self | href | |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

**Update the cluster-wide SNMP configuration**

PATCH `/support/snmp`

**Introduced In:** 9.7

Updates the cluster wide SNMP configuration, such as:

- enabling or disabling SNMP
- enabling or disabling SNMP traps
- enabling or disabling authentication traps
- setting the contact and location information for the SNMP server
- triggering an SNMP test trap

#### Related ONTAP commands

- `options snmp.enable`
- `system snmp authtrap`
- `system snmp init`

#### Learn more

- DOC /support/snmp

#### Request Body

| Name | Type | Description |
|---|---|---|
| auth_traps_enabled | boolean | Specifies whether to enable or disable SNMP authentication traps. |
| contact | string | Specifies the contact person for the SNMP server |
| enabled | boolean | Specifies whether to enable or disable SNMP. |
| location | string | Specifies the location of the SNMP server |
| traps_enabled | boolean | Specifies whether to enable or disable SNMP traps. |
| trigger_test_trap | boolean | Trigger a test SNMP trap. |

#### Example request

```
{
  "auth_traps_enabled": 1,
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "enabled": 1,
  "location": "Building 1",
  "traps_enabled": 1,
  "trigger_test_trap": 1
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 1179738 | Cluster identification not available for modification since a cluster has not been created. |
| 1179739 | Cluster identification modification is not available in maintenance mode or other boot modes. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

self_link

snmp

Cluster-wide SNMP configuration.

| Name | Type | Description |
|---|---|---|
| auth_traps_enabled | boolean | Specifies whether to enable or disable SNMP authentication traps. |
| contact | string | Specifies the contact person for the SNMP server |
| enabled | boolean | Specifies whether to enable or disable SNMP. |
| location | string | Specifies the location of the SNMP server |
| traps_enabled | boolean | Specifies whether to enable or disable SNMP traps. |
| trigger_test_trap | boolean | Trigger a test SNMP trap. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |

| Name | Type | Description |
| --- | --- | --- |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Create or retrieve SNMP traphosts

### Create or retrieve SNMP traphosts

#### Overview

This SNMP endpoint is used to retrieve, configure and reconfigure SNMP traphosts. An SNMP traphost is a host machine that is configured to receive SNMP traps from ONTAP. An SNMP traphost can be an SNMPv1/SNMPv2c traphost or an SNMPv3 traphost. ONTAP sends SNMPv1/SNMPv2c traps to an SNMPv1/SNMPv2c traphost over the SNMPv1/SNMPv2c protocol. Similarly, it sends SNMPv3 traps to an SNMPv3 traphost over the SNMPv3 protocol.

#### Examples

### Retrieves a list of SNMP traphosts in the cluster, sorted by IP address

```
# The API:
GET "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts?order_by=ip_address"

# The response:
{
"records": [
  {
    "host": "example_host_name",
    "ip_address": "12.34.56.78",
    "_links": {
      "self": {
        "href": "/api/support/snmp/traphosts/example_host_name"
      }
    }
  },
  {
    "host": "example_host_name2",
    "ip_address": "12.34.56.90",
    "_links": {
      "self": {
        "href": "/api/support/snmp/traphosts/example_host_name2"
      }
    }
  },
  {
    "host": "example_host_name3",
    "ip_address": "12.34.56.91",
    "_links": {
      "self": {
        "href": "/api/support/snmp/traphosts/example_host_name3"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/support/snmp/traphosts?order_by=ip_address"
  }
}
}
```

**Retrieves a list of SNMP traphosts in the cluster for which SNMP user is 'snmpv3user1'**

```
# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts?name=snmpv3user1"

# The response:
{
"records": [
  {
    "host": "example_host_name",
    "user": {
      "name": "snmpv3user1",
      "_links": {
        "self": {
          "href":
"/api/support/snmp/users/800003150558b57e8dbd9ce9119d82005056a7b4e5/snmpv3
user1"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/snmp/traphosts/example_host_name"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/support/snmp/traphosts?user.name=snmpv3user1"
  }
}
}
```

**Adds an SNMPv1 or SNMPv2c traphost**

```
# The API:
POST "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/traphosts" -d
'{"host":"10.140.128.48"}'

# The response:
201 Created
{}
```

#### Adds an SNMPv3 traphost

```
# The API:
POST "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/traphosts" -d '{"host":"fd-
geoycpi-uno.gycpi.b.yahoodns.net", "user":{"name":"snmpv3user1"}}'

# The response:
201 Created
{}
```

### Retrieve SNMP traphosts

GET `/support/snmp/traphosts`

**Introduced In:** 9.7

Retrieves the list of SNMP traphosts along with the SNMP users configured for those traphosts.

#### Related ONTAP commands

- `system snmp traphost show`

#### Learn more

- DOC /support/snmp/traphosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| host | string | query | False | Filter by host |
| user.name | string | query | False | Filter by user.name |
| ip_address | string | query | False | Filter by ip_address |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[snmp_traphost] | |

**Example response**


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "host": "traphost.example.com",
      "ip_address": "10.10.10.7",
      "user": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "snmpv3user3"
      }
    }
  ]
}
```


**Error**

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

href

| Name | Type | Description |
| --- | --- | --- |
| href | string | |

_links

| Name | Type | Description |
| --- | --- | --- |
| next | href | |
| self | href | |

_links

| Name | Type | Description |
| --- | --- | --- |
| self | href | |

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |

| Name | Type | Description |
| --- | --- | --- |
| name | string | Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.<br>• example: snmpv3user3<br>• Introduced in: 9.7<br>• x-nullable: true |

snmp_traphost

SNMP manager or host machine that receives SNMP traps from ONTAP.

| Name | Type | Description |
| --- | --- | --- |
| _links | _links | |
| host | string | Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |
| ip_address | string | IPv4 or IPv6 address |

| Name | Type | Description |
| --- | --- | --- |
| user | user | Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Create SNMP traphosts

POST `/support/snmp/traphosts`

**Introduced In:** 9.7

Creates SNMP traphosts. While adding an SNMPv3 traphost, an SNMPv3 user configured in ONTAP must be specified. ONTAP uses this user's credentials to authenticate and/or encrypt traps sent to this SNMPv3 traphost. While adding an SNMPv1/SNMPv2c traphost, SNMPv1/SNMPv2c user or community need not be specified.

#### Required properties

- `host` - Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost.

#### Recommended optional properties

- If `host` refers to an SNMPv3 traphost, the following field is required:

- `user` - SNMPv3 or User-based Security Model (USM) user.

- For an SNMPv1/SNMPv2c traphost, ONTAP automatically uses 'public' if 'public' is configured or no community is configured. Otherwise, ONTAP uses the first configured community.

#### Related ONTAP commands

- `system snmp traphost add`

#### Learn more

- DOC /support/snmp/traphosts

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |

#### Request Body

| Name | Type | Description |
|---|---|---|
| host | string | Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |
| user | user | Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community. |

### Example request

```
{
  "host": "traphost.example.com",
  "user": {
    "name": "snmpv3user3"
  }
}
```

#### Response

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| num_records | integer | Number of records. |
| records | array[snmp_traphost] | |

#### Example response

```
{
  "num_records": 1,
  "records": [
    {
      "host": "traphost.example.com",
      "user": {
        "name": "snmpv3user3"
      }
    }
  ]
}
```

#### Headers

| Name | Description | Type |
|---|---|---|
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8847365 | Unknown host. |
| 9043969 | Duplicate traphost entry. |
| 9043980 | IPv6 support is disabled. |
| 9043991 | Not a USM user. |

| Error Code | Description |
| --- | --- |
| 9043993 | Current cluster version does not support SNMPv3 traps. |
| 9044001 | Failed to create SNMPv1 traphost. |
| 9044002 | Failed to create SNMPv3 traphost. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

| Name | Type | Description |
|---|---|---|
| name | string | Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.<br>• example: snmpv3user3<br>• Introduced in: 9.7<br>• x-nullable: true |

snmp_traphost

SNMP manager or host machine that receives SNMP traps from ONTAP.

| Name | Type | Description |
|---|---|---|
| host | string | Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |
| user | user | Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage individual SNMP traphosts

### Manage individual SNMP traphosts

#### Overview

Defines, retrieves or deletes an individual SNMP traphost.

#### Examples

### Retrieves an individual traphost in the cluster

```
# The API:
GET "/api/support/snmp/traphosts/{host}"

# The call
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts/10.235.36.62"

# The response
{
"host": "example_host_name",
"ip_address": "10.235.36.62",
"user": {
  "name": "public",
  "_links": {
    "self": {
      "href":
"/api/support/snmp/users/800003150558b57e8dbd9ce9119d82005056a7b4e5/public
"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/support/snmp/traphosts/10.235.36.62"
  }
}
}
```

**Deletes an individual traphost in the cluster**

```
# The API:
DELETE "/api/support/snmp/traphosts/{host}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
DELETE "https://<mgmt-
ip>/api/support/snmp/traphosts/3ffe:ffff:100:f102::1"

# The response:
200 OK
{}
```

### Delete an SNMP traphost

DELETE `/support/snmp/traphosts/{host}`

**Introduced In:** 9.7

Deletes an SNMP traphost.

#### Learn more

- DOC /support/snmp/traphosts/{host}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| host | string | path | True | Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

**Error**

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8847365 | Unknown host. |

| Error Code | Description |
| --- | --- |
| 9043970 | Traphost entry does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

**Example error**

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

**Definitions**

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve SNMP traphost details

GET `/support/snmp/traphosts/{host}`

**Introduced In:** 9.7

Retrieves the details of an SNMP traphost along with the SNMP user configured for that traphost.

#### Learn more

- DOC /support/snmp/traphosts/{host}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| host | string | path | True | Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |
| fields | array[string] | query | False | Specify the fields to return. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| host | string | Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost. |
| ip_address | string | IPv4 or IPv6 address |
| user | user | Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community. |

#### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "host": "traphost.example.com",
  "ip_address": "10.10.10.7",
  "user": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "snmpv3user3"
  }
}
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 8847365 | Unknown host. |
| 9043970 | Traphost entry does not exist. |

Also see the table of common errors in the Response body overview section of this documentation.

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.<br>• example: snmpv3user3<br>• Introduced in: 9.7<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Retrieve and create SNMP users

### Retrieve and create SNMP users

#### Overview

Defines, retrieves and adds new SNMP users. An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. Use a "community" user to query ONTAP SNMP server over SNMPv1/SNMPv2c protocol. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. Use a local SNMPv3 user to query ONTAP SNMP server over SNMPv3 and/or to send SNMPv3 traps. A remote SNMPv3 user is configured in ONTAP as well as on a remote switch. ONTAP SNMP client uses a remote SNMPv3 user to query the switch over SNMPv3.

#### Examples

#### Retrieves a list of SNMP users in the cluster, sorted by SVM UUID

```
# The API:
GET "/api/support/snmp/users"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/users?order_by=owner.uuid"

# The response:
{
```

```
"records": [
  {
    "engine_id": "80000315052d7d2c4410b8e911bc9f005056bb942e",
    "name": "snmpv1user1",
    "owner": {
      "uuid": "442c7d2d-b810-11e9-bc9f-005056bb942e"
    },
    "_links": {
      "self": {
        "href":
"/api/support/snmp/users/80000315052d7d2c4410b8e911bc9f005056bb942e/snmpv1
user1"
      }
    }
  },
  {
    "engine_id": "80000315058e02057c0fb8e911bc9f005056bb942e",
    "name": "snmpv3user1",
    "owner": {
      "uuid": "7c05028e-b80f-11e9-bc9f-005056bb942e"
    },
    "_links": {
      "self": {
        "href":
"/api/support/snmp/users/80000315058e02057c0fb8e911bc9f005056bb942e/snmpv3
user1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/support/snmp/users?order_by=owner.uuid"
  }
}
}
```


**Creates an SNMPv1 or SNMPv2c community user**

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"}, "name":"public", "authentication_method":"community"}'

# The response:
201 Created
{}
```

**Creates a local SNMPv3 user**

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"}, "name":"snmpv3u1", "authentication_method":"usm",
"snmpv3":{"authentication_protocol":"sha",
"authentication_password":"jelly22fi$h", "privacy_protocol":"aes",
"privacy_password":"a11Black$"}}'

# The response:
201 Created
{}
```

**Creates a remote SNMPv3 user**

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"},
"engine_id":"8000031505b67667a26975e9118a480050568e6f74",
"name":"switchuser1", "authentication_method":"usm",
"switch_address":"fdf8:f53b:82e4::53",
"snmpv3":{"authentication_protocol":"md5",
"authentication_password":"TfhIeliw601lS.Rw$3pm"}}'

# The response:
201 Created
{}
```

### Retrieve SNMP users on a cluster

GET `/support/snmp/users`

**Introduced In:** 9.7

Retrieves the list of SNMP users on the cluster.

#### Related ONTAP commands

- `security snmpusers`
- `security login show -application snmp`

#### Learn more

- DOC /support/snmp/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engine_id | string | query | False | Filter by engine_id |
| authentication_method | string | query | False | Filter by authentication_method |
| scope | string | query | False | Filter by scope |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| owner.name | string | query | False | Filter by owner.name |
| owner.uuid | string | query | False | Filter by owner.uuid |
| snmpv3.privacy_protocol | string | query | False | Filter by snmpv3.privacy_protocol |
| snmpv3.authentication_protocol | string | query | False | Filter by snmpv3.authentication_protocol |
| comment | string | query | False | Filter by comment<br>• maxLength: 128<br>• minLength: 0 |
| switch_address | string | query | False | Filter by switch_address |
| name | string | query | False | Filter by name<br>• maxLength: 32 |
| fields | array[string] | query | False | Specify the fields to return. |
| max_records | integer | query | False | Limit the number of records returned. |
| return_records | boolean | query | False | The default is true for GET calls. When set to false, only the number of records is returned.<br>• Default value: 1 |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |
| order_by | array[string] | query | False | Order results by specified fields and optional [asc |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| num_records | integer | Number of records. |
| records | array[snmp_user] | |

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_method": "usm",
      "comment": "This is a comment.",
      "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
      "name": "snmpv3user2",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "svm",
      "snmpv3": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "authentication_protocol": "sha2_256",
        "privacy_protocol": "aes128"
      },
      "switch_address": "10.23.34.45"
    }
  ]
}
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| next | href | |
| self | href | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_protocol | string | Authentication protocol. |
| privacy_protocol | string | Privacy protocol. |

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_method | string | Optional authentication method. |
| comment | string | Optional comment text. |
| engine_id | string | Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier. |
| name | string | SNMP user name. |
| owner | owner | Optional name and UUID of owning Storage Virtual Machine (SVM). |
| scope | string | Set to "svm" for data Storage Virtual Machine (SVM) SNMP users and to "cluster" for administrative SVM SNMP users. |

| Name | Type | Description |
| --- | --- | --- |
| snmpv3 | snmpv3 | Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user. |
| switch_address | string | Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.<br>• example: 10.23.34.45<br>• Introduced in: 9.7<br>• readCreate: 1<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |

| Name | Type | Description |
| --- | --- | --- |
| target | string | The target parameter that caused the error. |

### Create a cluster-scoped or an SVM-scoped SNMP user

POST `/support/snmp/users`

**Introduced In:** 9.7

Creates either a cluster-scoped or an SVM-scoped SNMP user. This user can be an SNMPv1 or SNMPv2c community user or an SNMPv3 user. An SNMPv3 user can be a local SNMPv3 user or a remote SNMPv3 user.

#### Required properties

- `name` - SNMP user name

#### Optional properties

- `owner` - Name and UUID of owning SVM.
- `engine_id` - Engine ID of owning SVM or remote switch.
- `authentication_method` - Authentication method
- `switch_address` - Optional remote switch address
- `snmpv3` - SNMPv3-specific credentials
- `comment` - Comment text

#### Default property values

- `snmpv3.authentication_protocol` - none
- `snmpv3.privacy_protocol` - none

#### Related ONTAP commands

- `security login create`
- `system snmp community add`

#### Learn more

- DOC /support/snmp/users

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_records | boolean | query | False | The default is false. If set to true, the records are returned.<br>• Default value: |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| authentication_method | string | Optional authentication method. |
| comment | string | Optional comment text. |

| Name | Type | Description |
| --- | --- | --- |
| engine_id | string | Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier. |
| name | string | SNMP user name. |
| owner | owner | Optional name and UUID of owning Storage Virtual Machine (SVM). |
| snmpv3 | snmpv3 | Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user. |

| Name | Type | Description |
| --- | --- | --- |
| switch_address | string | Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.<br>• example: 10.23.34.45<br>• Introduced in: 9.7<br>• readCreate: 1<br>• x-nullable: true |

**Example request**

```
{
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "snmpv3": {
    "authentication_password": "humTdumt*@t0nAwa11",
    "authentication_protocol": "sha2_256",
    "privacy_password": "p@**GOandCLCt*200",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```

**Response**

```
Status: 201, Created
```

| Name | Type | Description |
|---|---|---|
| authentication_method | string | Optional authentication method. |
| comment | string | Optional comment text. |
| engine_id | string | Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier. |
| name | string | SNMP user name. |
| owner | owner | Optional name and UUID of owning Storage Virtual Machine (SVM). |
| snmpv3 | snmpv3 | Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user. |

| Name | Type | Description |
| --- | --- | --- |
| switch_address | string | Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.<br>• example: 10.23.34.45<br>• Introduced in: 9.7<br>• readCreate: 1<br>• x-nullable: true |

### Example response


```
{
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "snmpv3": {
    "authentication_password": "humTdumt*@t0nAwa11",
    "authentication_protocol": "sha2_256",
    "privacy_password": "p@**GOandCLCt*200",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```


### Headers

| Name | Description | Type |
| --- | --- | --- |
| Location | Useful for tracking the resource location | string |

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
| --- | --- |
| 1263341 | Invalid authentication protocol. |
| 1263342 | Invalid privacy protocol. |
| 1263567 | Authentication password not supplied. |
| 1263568 | Privacy password not supplied. |
| 1263569 | Authentication password does not conform to ONTAP requirements. |
| 1263570 | Privacy password does not conform to ONTAP requirements. |
| 2621475 | This operation is not allowed on a node SVM. |
| 2621601 | This operation is not permitted on a system SVM. |
| 5636101 | One role for all applications and authentication methods is recommended for a user account. |
| 5636106 | Engine ID starting with first bit set as 1 in first octet should have a minimum of 5 or maximum of 32 octets. |
| 5636107 | Engine ID specified is based on an IP address, and must therefore be 9 octets in length. |
| 5636108 | Engine ID specified is based on an IPv6 address, and must therefore be 21 octets in length. |
| 5636109 | Engine ID specified is based on a MAC address, and must therefore be 11 octets in length. |
| 5636110 | Engine ID starting with first bit set as 0 in the first octet needs to be 12 octets in length. |
| 5636111 | Engine ID cannot have all zeroes. |
| 5636112 | Invalid Engine ID specified. |
| 5636113 | Cannot enable SNMP user capability for data SVM. |
| 5636116 | Cannot read SNMP user capability for data SVM. |
| 5636121 | User account name is reserved for use by the system. |

| Error Code | Description |
| --- | --- |
| 5636123 | Cannot create an SNMP user with a role other than readonly, none, or admin. |
| 5636124 | Cannot create an SNMP user with a role other than vsadmin-readonly, none, or vsadmin. |
| 5636148 | Failed to discover SNMPv3 engine ID of remote SNMP agent. |
| 5636150 | The switch_address parameter can be specified only for administrative SVMs. |
| 5636152 | 0.0.0.0 cannot be specified as the IP Address for the switch_address parameter. |
| 5636153 | Engine ID is already associated with user. |
| 7077897 | Invalid character in username. Valid characters":" A-Z, a-z, 0-9, ".", "_", and "-". Note that the user name cannot start with "-". |
| 7077899 | Username does not meet length requirements. |
| 7077906 | Cannot use given role with this SVM because a role with that name has not been defined for the SVM. |
| 9043995 | Failed to create SNMP community. Reason":" SNMPv1 and SNMPv2c are not supported when cluster FIPS mode is enabled. |
| 9043999 | ONTAP failed to create an SNMPv3 user because SNMPv3 is disabled in the cluster. |
| 9044020 | Customized engine-id cannot be specified for community users. |
| 9044021 | The specified authentication protocol is not supported when cluster FIPS mode is enabled. |
| 9044022 | The specified privacy protocol is not supported when cluster FIPS mode is enabled. |
| 9044023 | The switch_address parameter can be specified only for creating remote SNMPv3 users. |

Also see the table of common errors in the Response body overview section of this documentation.

**Definitions**

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

| Name | Type | Description |
|---|---|---|
| authentication_password | string | Authentication protocol password. |
| authentication_protocol | string | Authentication protocol. |
| privacy_password | string | Privacy protocol password. |
| privacy_protocol | string | Privacy protocol. |

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

| Name | Type | Description |
| --- | --- | --- |
| authentication_method | string | Optional authentication method. |
| comment | string | Optional comment text. |
| engine_id | string | Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier. |
| name | string | SNMP user name. |
| owner | owner | Optional name and UUID of owning Storage Virtual Machine (SVM). |
| snmpv3 | snmpv3 | Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user. |

| Name | Type | Description |
| --- | --- | --- |
| switch_address | string | Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.<br>• example: 10.23.34.45<br>• Introduced in: 9.7<br>• readCreate: 1<br>• x-nullable: true |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

## Manage individual SNMP users

## Manage individual SNMP users

### Overview

Defines, retrieves, updates and deletes an individual SNMP user.

### Examples

### Retrieves the details of an SNMP user

```
# The API:
GET "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/users/80000315056622e52625a9e911a981005056bb1dcb/snmp
v1user2"

# The response:
{
"engine_id": "80000315056622e52625a9e911a981005056bb1dcb",
"name": "snmpv1user2",
"scope": "cluster",
"owner": {
  "name": "cluster-1",
  "uuid": "26e52266-a925-11e9-a981-005056bb1dcb"
},
"authentication_method": "community",
"_links": {
  "self": {
    "href":
"/api/support/snmp/users/80000315056622e52625a9e911a981005056bb1dcb/snmpv1
user2"
  }
}
}
```

### Updates the comment parameter for an individual SNMP user

```
# The API:
PATCH "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-
ip>/api/support/snmp/users/8000031505b67667a26975e9118a480050568e6f74/publ
ic" -d '{"comment":"Default SNMP community"}'

# The response:
200 OK
{}
```

#### Deletes an individual SNMP user in the cluster

```
# The API:
DELETE "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
DELETE "https://<mgmt-
ip>/api/support/snmp/users/8000031505b67667a26975e9118a480050568e6f74/snmp
user"

# The response:
200 OK
{}
```

### Delete an SNMP user

DELETE `/support/snmp/users/{engine_id}/{name}`

**Introduced In:** 9.7

Deletes an SNMP user. The engine ID can be the engine ID of the administrative SVM or a data SVM. It can also be the SNMPv3 engine ID of a remote switch.

#### Related ONTAP commands

- `security login delete`
- `system snmp community delete`

#### Learn more

- DOC /support/snmp/users/{engine_id}/{name}

**Parameters**

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engine_id | string | path | True | Engine ID of owning SVM or remote switch. |
| name | string | path | True | SNMP user name. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

| Name | Type | Description |
| --- | --- | --- |
| error | returned_error | |

#### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

#### Definitions

**See Definitions**

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Retrieve SNMP user details

GET `/support/snmp/users/{engine_id}/{name}`

**Introduced In:** 9.7

Retrieves the details of an SNMP user. The engine ID can be the engine ID of the administrative SVM or a data SVM. It can also be the SNMPv3 engine ID of a remote switch.

#### Related ONTAP commands

- `security snmpusers -vserver {SVM name} -username <User Name>`
- `security login show -application snmp -vserver {SVM name} -user-or-group-name <User Name>`

#### Learn more

- DOC /support/snmp/users/{engine_id}/{name}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engine_id | string | path | True | Engine ID of owning SVM or remote switch. |
| name | string | path | True | SNMP user name. |
| fields | array[string] | query | False | Specify the fields to return. |
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.<br>• Default value: 15<br>• Max value: 120<br>• Min value: 0 |

**Response**

```
Status: 200, Ok
```

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_method | string | Optional authentication method. |
| comment | string | Optional comment text. |

| Name | Type | Description |
| --- | --- | --- |
| engine_id | string | Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier. |
| name | string | SNMP user name. |
| owner | owner | Optional name and UUID of owning Storage Virtual Machine (SVM). |
| scope | string | Set to "svm" for data Storage Virtual Machine (SVM) SNMP users and to "cluster" for administrative SVM SNMP users. |
| snmpv3 | snmpv3 | Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user. |

| Name | Type | Description |
| --- | --- | --- |
| switch_address | string | Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.<br>• example: 10.23.34.45<br>• Introduced in: 9.7<br>• readCreate: 1<br>• x-nullable: true |

#### Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "svm",
  "snmpv3": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_protocol": "sha2_256",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```


#### Error

```
Status: Default, Error
```

| Name | Type | Description |
|---|---|---|
| error | returned_error | |

### Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

| Name | Type | Description |
|---|---|---|
| self | href | |

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

| Name | Type | Description |
|---|---|---|
| _links | _links | |
| authentication_protocol | string | Authentication protocol. |
| privacy_protocol | string | Privacy protocol. |

error_arguments

| Name | Type | Description |
|---|---|---|
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
|---|---|---|
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

### Update the comment parameter of an SNMP user

PATCH `/support/snmp/users/{engine_id}/{name}`

**Introduced In:** 9.7

Updates the comment parameter of an SNMP user.

#### Optional properties

- `comment` - Comment text.

#### Related ONTAP commands

- `security login modify`

#### Learn more

- DOC /support/snmp/users/{engine_id}/{name}

#### Parameters

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| engine_id | string | path | True | Engine ID of owning SVM or remote switch. |
| name | string | path | True | SNMP user name. |

| Name | Type | In | Required | Description |
|---|---|---|---|---|
| return_timeout | integer | query | False | The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.<br>• Default value: 0<br>• Max value: 120<br>• Min value: 0 |

**Request Body**

| Name | Type | Description |
|---|---|---|
| comment | string | Optional comment text. |

**Example request**

```
{
  "comment": "This is a comment."
}
```

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

ONTAP Error Response Codes

| Error Code | Description |
|---|---|
| 2621475 | This operation is not allowed on a node SVM. |
| 2621699 | This operation is not allowed on a system SVM. |
| 5636123 | Cannot create an SNMP user with a role other than readonly, none, or admin. |
| 5636124 | Cannot create an SNMP user with a role other than vsadmin-readonly, none, or vsadmin. |
| 5832712 | Cannot modify attributes for user "diag." |
| 7077906 | Cannot use given role with this SVM because a role with that name has not been defined for the SVM. |
| 9043999 | ONTAP failed to create an SNMPv3 user because SNMPv3 is disabled in the cluster. |

Also see the table of common errors in the Response body overview section of this documentation.

#### Definitions

**See Definitions**

href

| Name | Type | Description |
|---|---|---|
| href | string | |

_links

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

| Name | Type | Description |
|---|---|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

| Name | Type | Description |
|---|---|---|
| authentication_protocol | string | Authentication protocol. |
| privacy_protocol | string | Privacy protocol. |

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

| Name | Type | Description |
|---|---|---|
| comment | string | Optional comment text. |

error_arguments

| Name | Type | Description |
| --- | --- | --- |
| code | string | Argument code |
| message | string | Message argument |

returned_error

| Name | Type | Description |
| --- | --- | --- |
| arguments | array[error_arguments] | Message arguments |
| code | string | Error code |
| message | string | Error message |
| target | string | The target parameter that caused the error. |

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

https://www.netapp.com/company/legal/copyright/

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

## Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

## Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

## API license

This documentation does not grant any license or other rights to the ONTAP REST API. Any use or access of the ONTAP REST API requires and is subject to a separate license agreement with NetApp.

## Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

**Copyright information**

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

**Trademark information**

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.